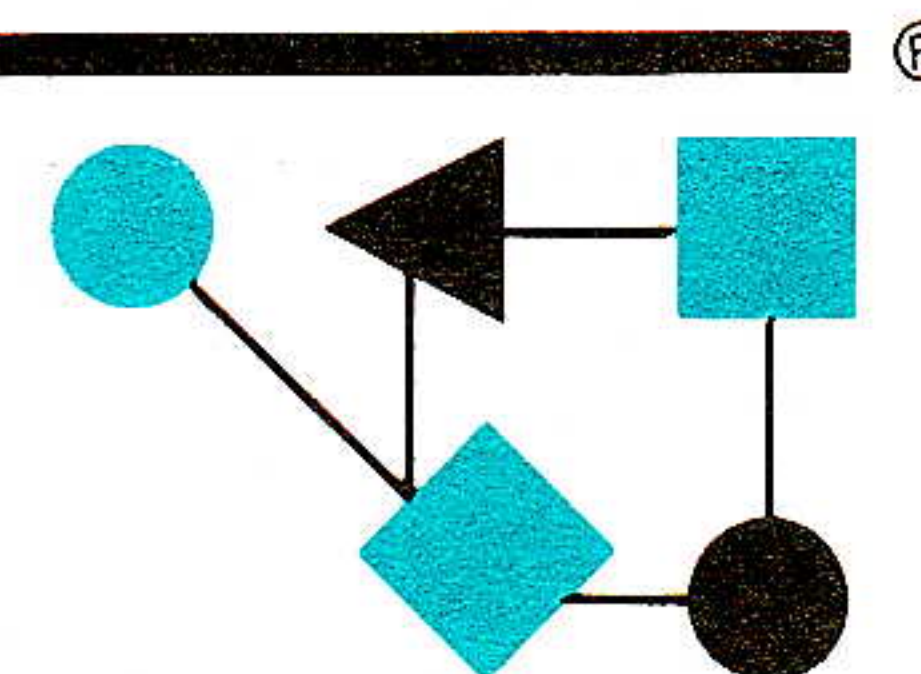


CONNEXIONS



The Interoperability Report

September 1991

Special Issue: Looking ahead to INTEROP 91 Fall

Volume 5, No. 9

*ConneXions —
The Interoperability Report
tracks current and emerging
standards and technologies
within the computer and
communications industry.*

In this issue:

Networks and Objects.....	2
Enterprise Networking.....	18
Networking the Telecom Standards Bodies.....	26
The INTEROP 90 network...	36
FDDI at INTEROP 91.....	40
Letters to the Editor.....	43
Announcements.....	44

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779. Toll-free: 1-800-INTEROP.

Copyright © 1991 by Interop, Inc.
Quotation with attribution encouraged.

ConneXions—The Interoperability Report
and the *ConneXions* logo are registered
trademarks of Interop, Inc.

ISSN 0894-5926

From the Editor

Every year, at the INTEROP conference and exhibition, we distribute a special issue of *ConneXions* with articles related to the conference sessions and other activities. This year is no exception, but in order to keep the INTEROP issue to a reasonable size we “start early” this month with a look at a few of the topics you’ll hear about at INTEROP 91 Fall.

We begin by examining object-oriented approaches to distributed computing. This important emerging technology is certain to have a great impact on the way in which computer networks of the future will operate. The article is by Michael Millikin of Gunstock Hill Associates. Mike will also be chairing a session on this topic at INTEROP 91 Fall.

Our conference has historically covered technology from several perspectives, mainly designer, implementor and user. This year we will continue this tradition with several *network case studies*. One of these is presented in this month’s issue. Wayne Russell describes the design and implementation of a large multivendor enterprise network at a government agency.

Traditionally, international telecommunications standards bodies such as CCITT and ISO have been slow in adopting networking technology for the dissemination of standards documents and other information. Anthony M. Rutkowski, counselor to the Secretary-General of the International Telecommunications Union, explains how this trend may soon change. You’ll hear much more about this in a session entitled “Who Owns Standards?” at INTEROP 91 Fall.

We also take a look *back*, to INTEROP 90, and go behind the scenes with one of the designers of the annual show network. Stev Knowles tells us what it is like to construct a complex network in a matter of a few hours. This effort will be undertaken again in a few weeks, as the shownet engineers build the INTEROP 91 network.

Finally, Mark Wolter gives an overview of tutorials, conference sessions, Solutions Showcase Demonstrations and other activities related to FDDI at this year’s show.

Coming next month, an in-depth look at LAN/WAN interconnect options, developments in SMDS, applications and techniques for LAN monitoring, background articles on the Great IGP Debate, and much more. Don’t miss the October issue and don’t miss INTEROP 91 Fall!

Networks and Objects

by Michael D. Millikin, Gunstock Hill Associates

Introduction

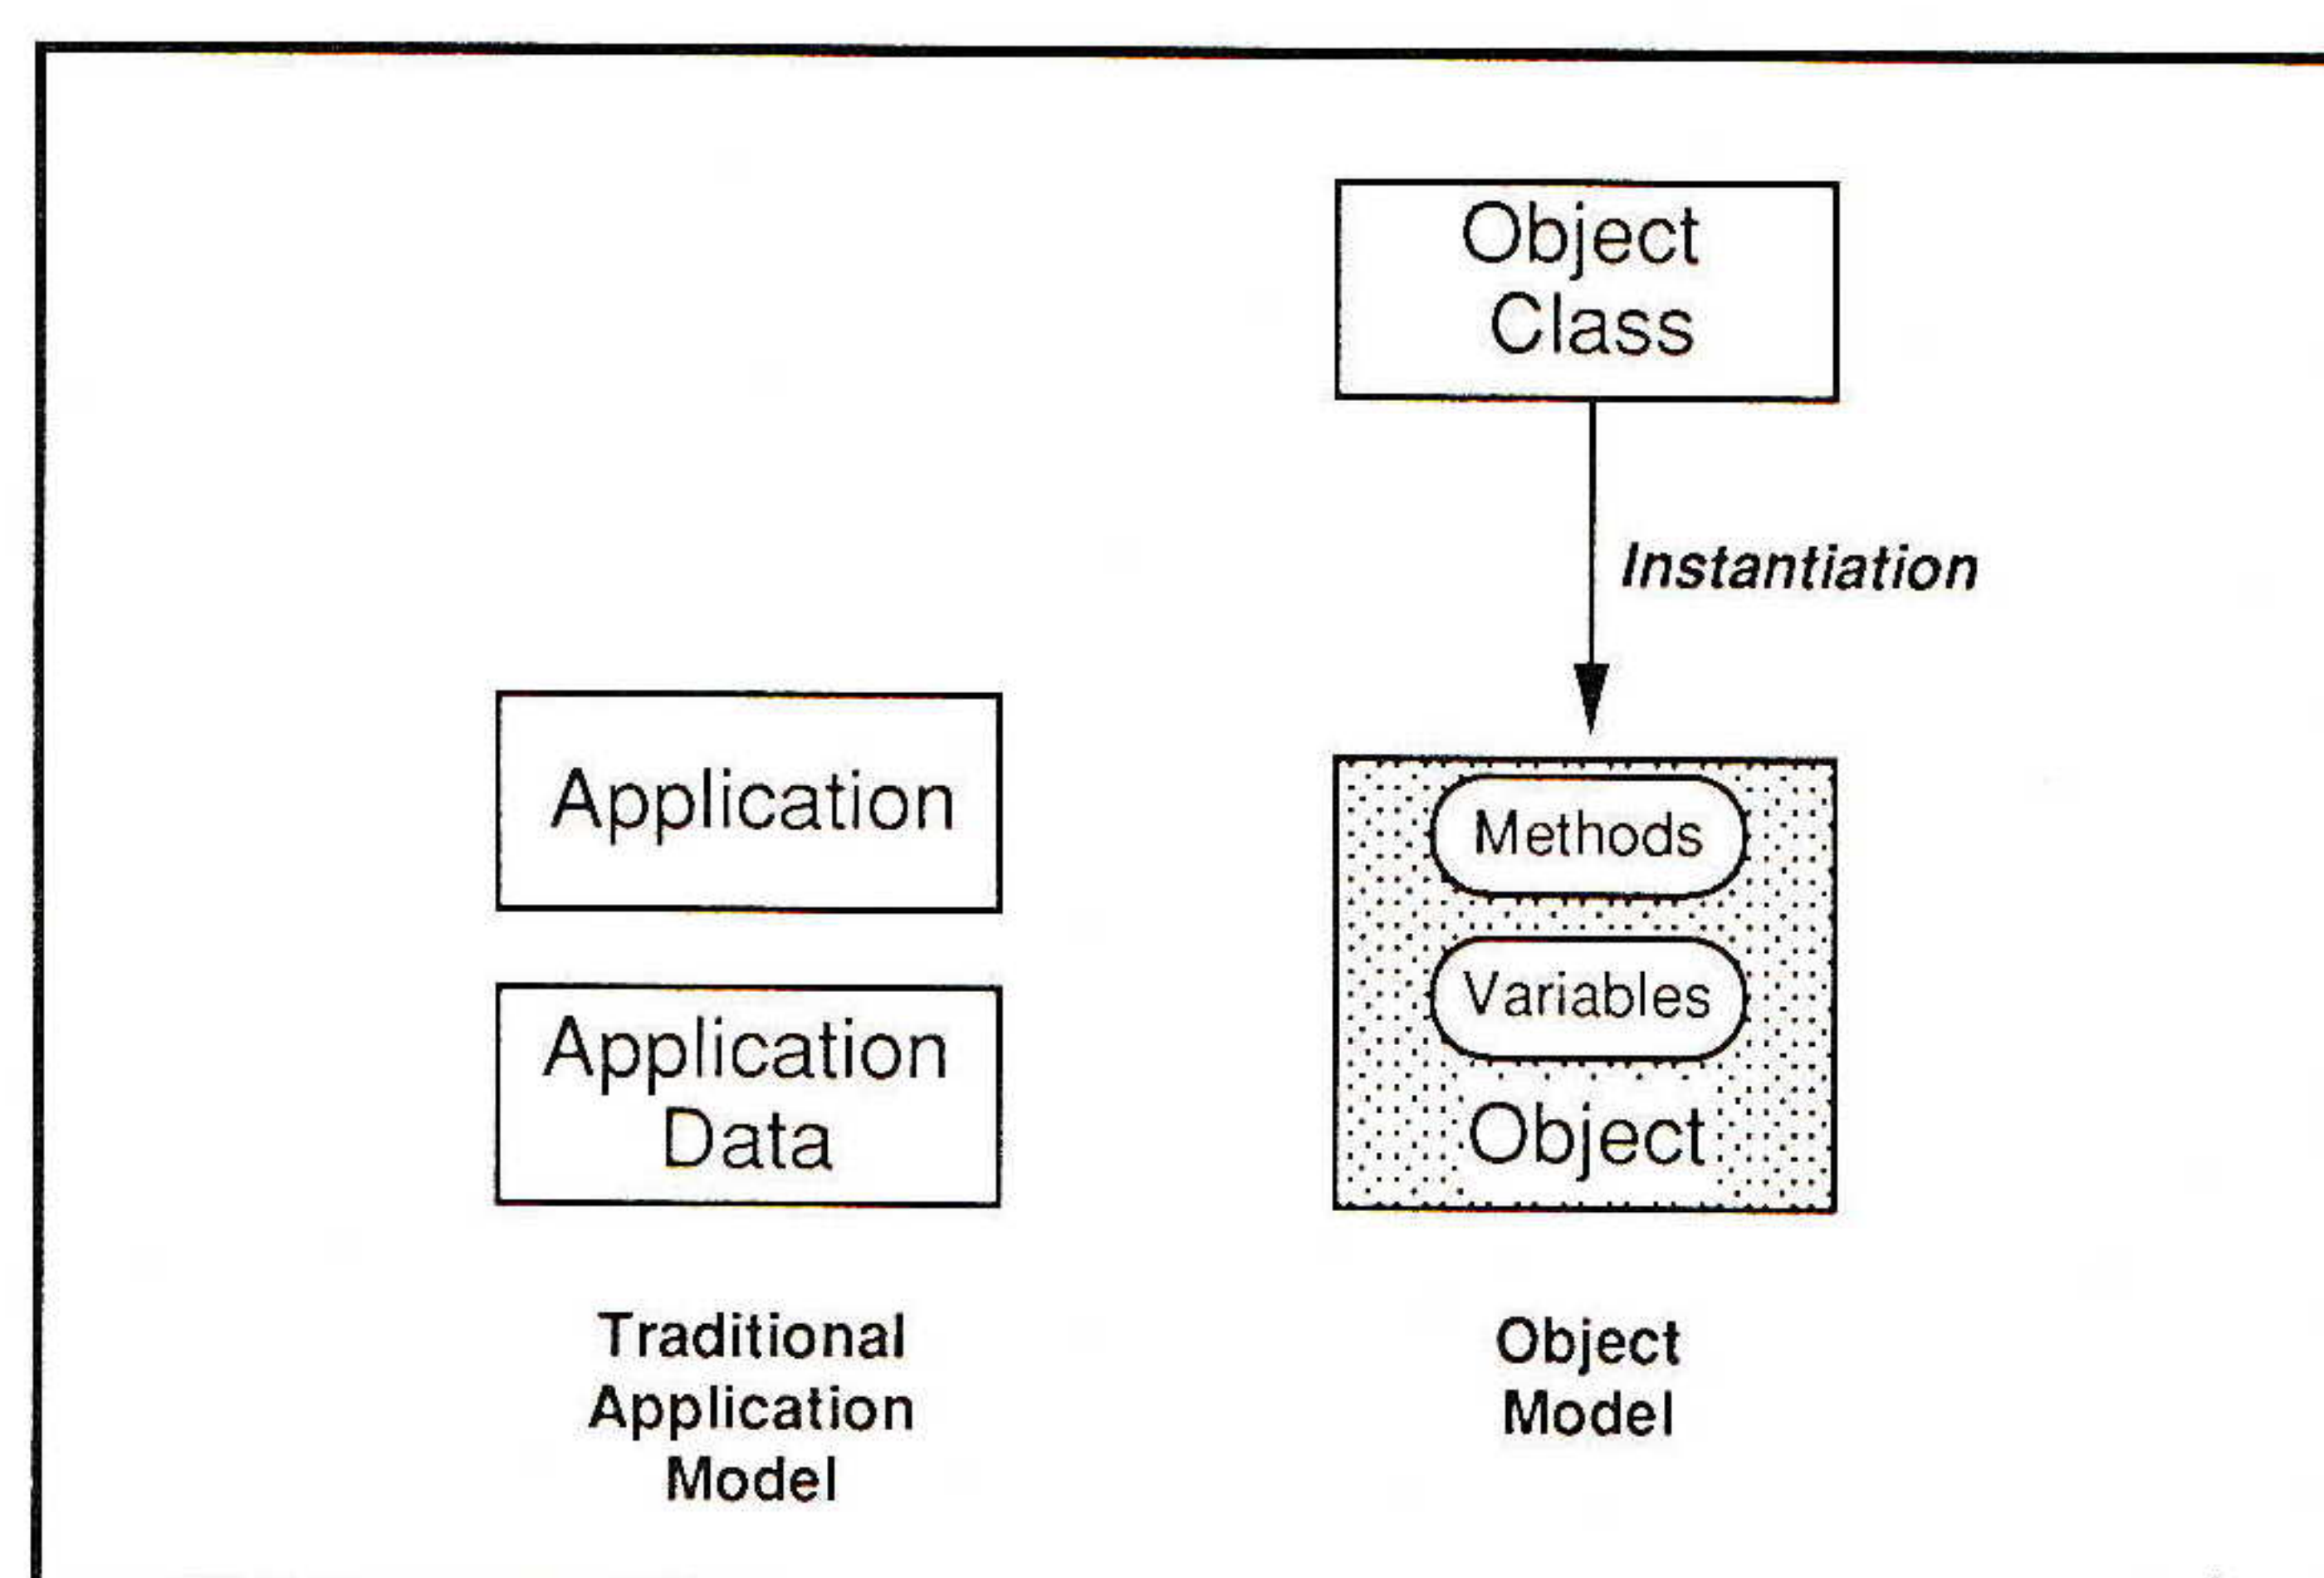
The high-level promise of the distributed computing environment is the ability to treat a heterogeneous, multivendor network as a single, logical system. Much of the current work in the industry, represented by the OSF *DCE* and the Unix International *Atlas*, is concerned with laying the infrastructure for those distributed environments that will carry us through the rest of the decade.

While that work is quite important, it is insufficient on its own to spur widespread adoption of the technology. The services are still too low-level, there are too many services to manage, and there is little uniformity offered above communications. To put it another way, the industry has to begin moving its focus for interoperability up the stack to a higher level. Object technology provides that necessary higher-level component to make a distributed system feasible and useful.

Terminology

Before proceeding any further, let's define our terms. A *distributed computing environment* is an architectural model of computing in which processes, resources, services and computers distributed across an heterogeneous network work together to perform their tasks in a manner transparent to users or other components of the network.

That's pretty straightforward. Defining *objects* is a little tougher, because objects can be, well, lots of things. An object is an instance of a type or class, and consists of variables and a set of methods that determine its behavior. Objects also have granularity, or size. Another, perhaps more immediately useful way, of thinking of objects is to consider a large-grained object comparable to an application and a particular set of data that works with that application encapsulated together. This would create objects at a file-level granularity. An example of this would be the objects in the NewWave environment. In turn, a small-grained object could be thought of as data encapsulated with a subroutine.



Where the current application model treats the application code and the data files as separate entities, the object model does away with such distinctions. Objects are self-describing entities. Each object contains the data (variables) and the operations (methods) required.

Figure 1: The object model versus the traditional model.

Even on a standalone workstation, objects need an *object manager*. The services such an object provides include:

- Controlling the life cycle of the object. The life cycle consists of actions such as creation, deletion, moving, copying and inserting.
- Activation of objects.
- Linking of objects.
- Persistent storage of objects.

In a distributed environment, there is a corresponding need for a distributed object management facility. This manager of distributed objects needs to be able to:

- Locate objects across the network.
- Operate across different network protocols.
- Send messages to other objects.

It must also be able to work well with existing applications (that is, applications that are procedural-based rather than tools that are object-based), and work with different persistent storage models (such as object databases, relational databases, and file systems). Additionally, it must be able to work with different types of object managers.

Objects offer well-defined operations and interfaces to other objects. The underlying storage mechanisms and algorithms are hidden. A consistent interface allows the replacement of underlying implementations. Standard object services support the integration of different types of objects from different sources. Objects can function as building blocks for special solutions or configurations through the reuse of existing object components as well as the addition of new ones. In short, objects provide a framework for managing complexity.

Neither distributed computing environments or object-oriented technology are new. In various forms, they have existed for more than a decade, and even longer in some cases. What is different, and what causes a great deal of heartburn to developers and implementors, is the application of such advanced technologies to the everyday business environment.

Heterogeneity and size

Two characteristics of this environment—heterogeneity and size—cause particular problems. Commercial information systems consist of a variety of hardware from different vendors running different operating systems. The distributed computing environment technologies (DCE, Atlas, et al.) are being designed to take care of providing the links to a certain level. But on top of those systems there are different types of applications, based on both object-oriented as well as procedural technology, working with different persistent storage models. A distributed object environment that hopes to garner any commercial acceptance must be able to work with that entire gamut of technology.

The issue of scale is simply that the distributed, object-based environments must be able to be implemented across the entire enterprise. The first requirement in meeting this need is in designing a system solution that can eventually handle hundreds of thousands of nodes and resources without sacrificing reliability or performance.

Networks and Objects (*continued*)

The second requirement—and possibly the more daunting—is finding user organizations willing to make the commitment to deploy such technology as the fundamental IS solution for the entire business. For commercial systems, this will mean a need to support transaction processing, as well as provide a comprehensive set of management and administration utilities.

These requirements put a definite burden on systems vendors that has caused some surprising, but necessary, partnering: such as that of Sun and HP with a Distributed Object Management Facility. The technology that we end up deploying must by design be capable of supporting multiple systems as well as multiple network protocols. Otherwise, that deployment just won't occur. No vendor is going to gain a strategic advantage in today's market climate by developing and touting an advanced technology that is proprietary.

However, neither technology is ready for prime time—yet. Some organizations, such as Citibank, are experimenting with distributed computing pilots, and are coding in those pilots with object-oriented languages. That is still far from the vision of the distributed, object-based system that will be the ultimate goal. Of the two technologies, distributed computing is closer to implementation than object orientation. However, distributed computing *needs* object orientation.

The payoff of a distributed environment as complex as that of the OSF's will not really be seen until it is deployed across the enterprise. That enterprise-wide scale is one of the design criteria that shaped the selection of the OSF technology. But deploying that technology on such a scale is a major technology transition for most corporations. Those businesses will not make that complete transition without a compelling business justification. Applications and solutions are the building blocks of such a justification. A highway system is no good without vehicles, and vehicles are pointless without drivers and cargo. Object technology will be providing those required solutions.

In this article, we are going to concentrate on object technology, but in the context of the distributed computing environment. We'll approach the issue from three points of view:

- First, using object technology to link existing types of applications
- Second, using object technology to develop new applications
- Third, interfaces between distributed objects.

What's in it for me?

Is this worth the trouble? If some companies are facing a backlash from entrenched MIS just against "open systems," what happens when we try to convince them about the need for a wholesale shift of their technology base?

Consider some of the following as examples of the possibilities of a distributed, object-based environment:

Compound documents/objects

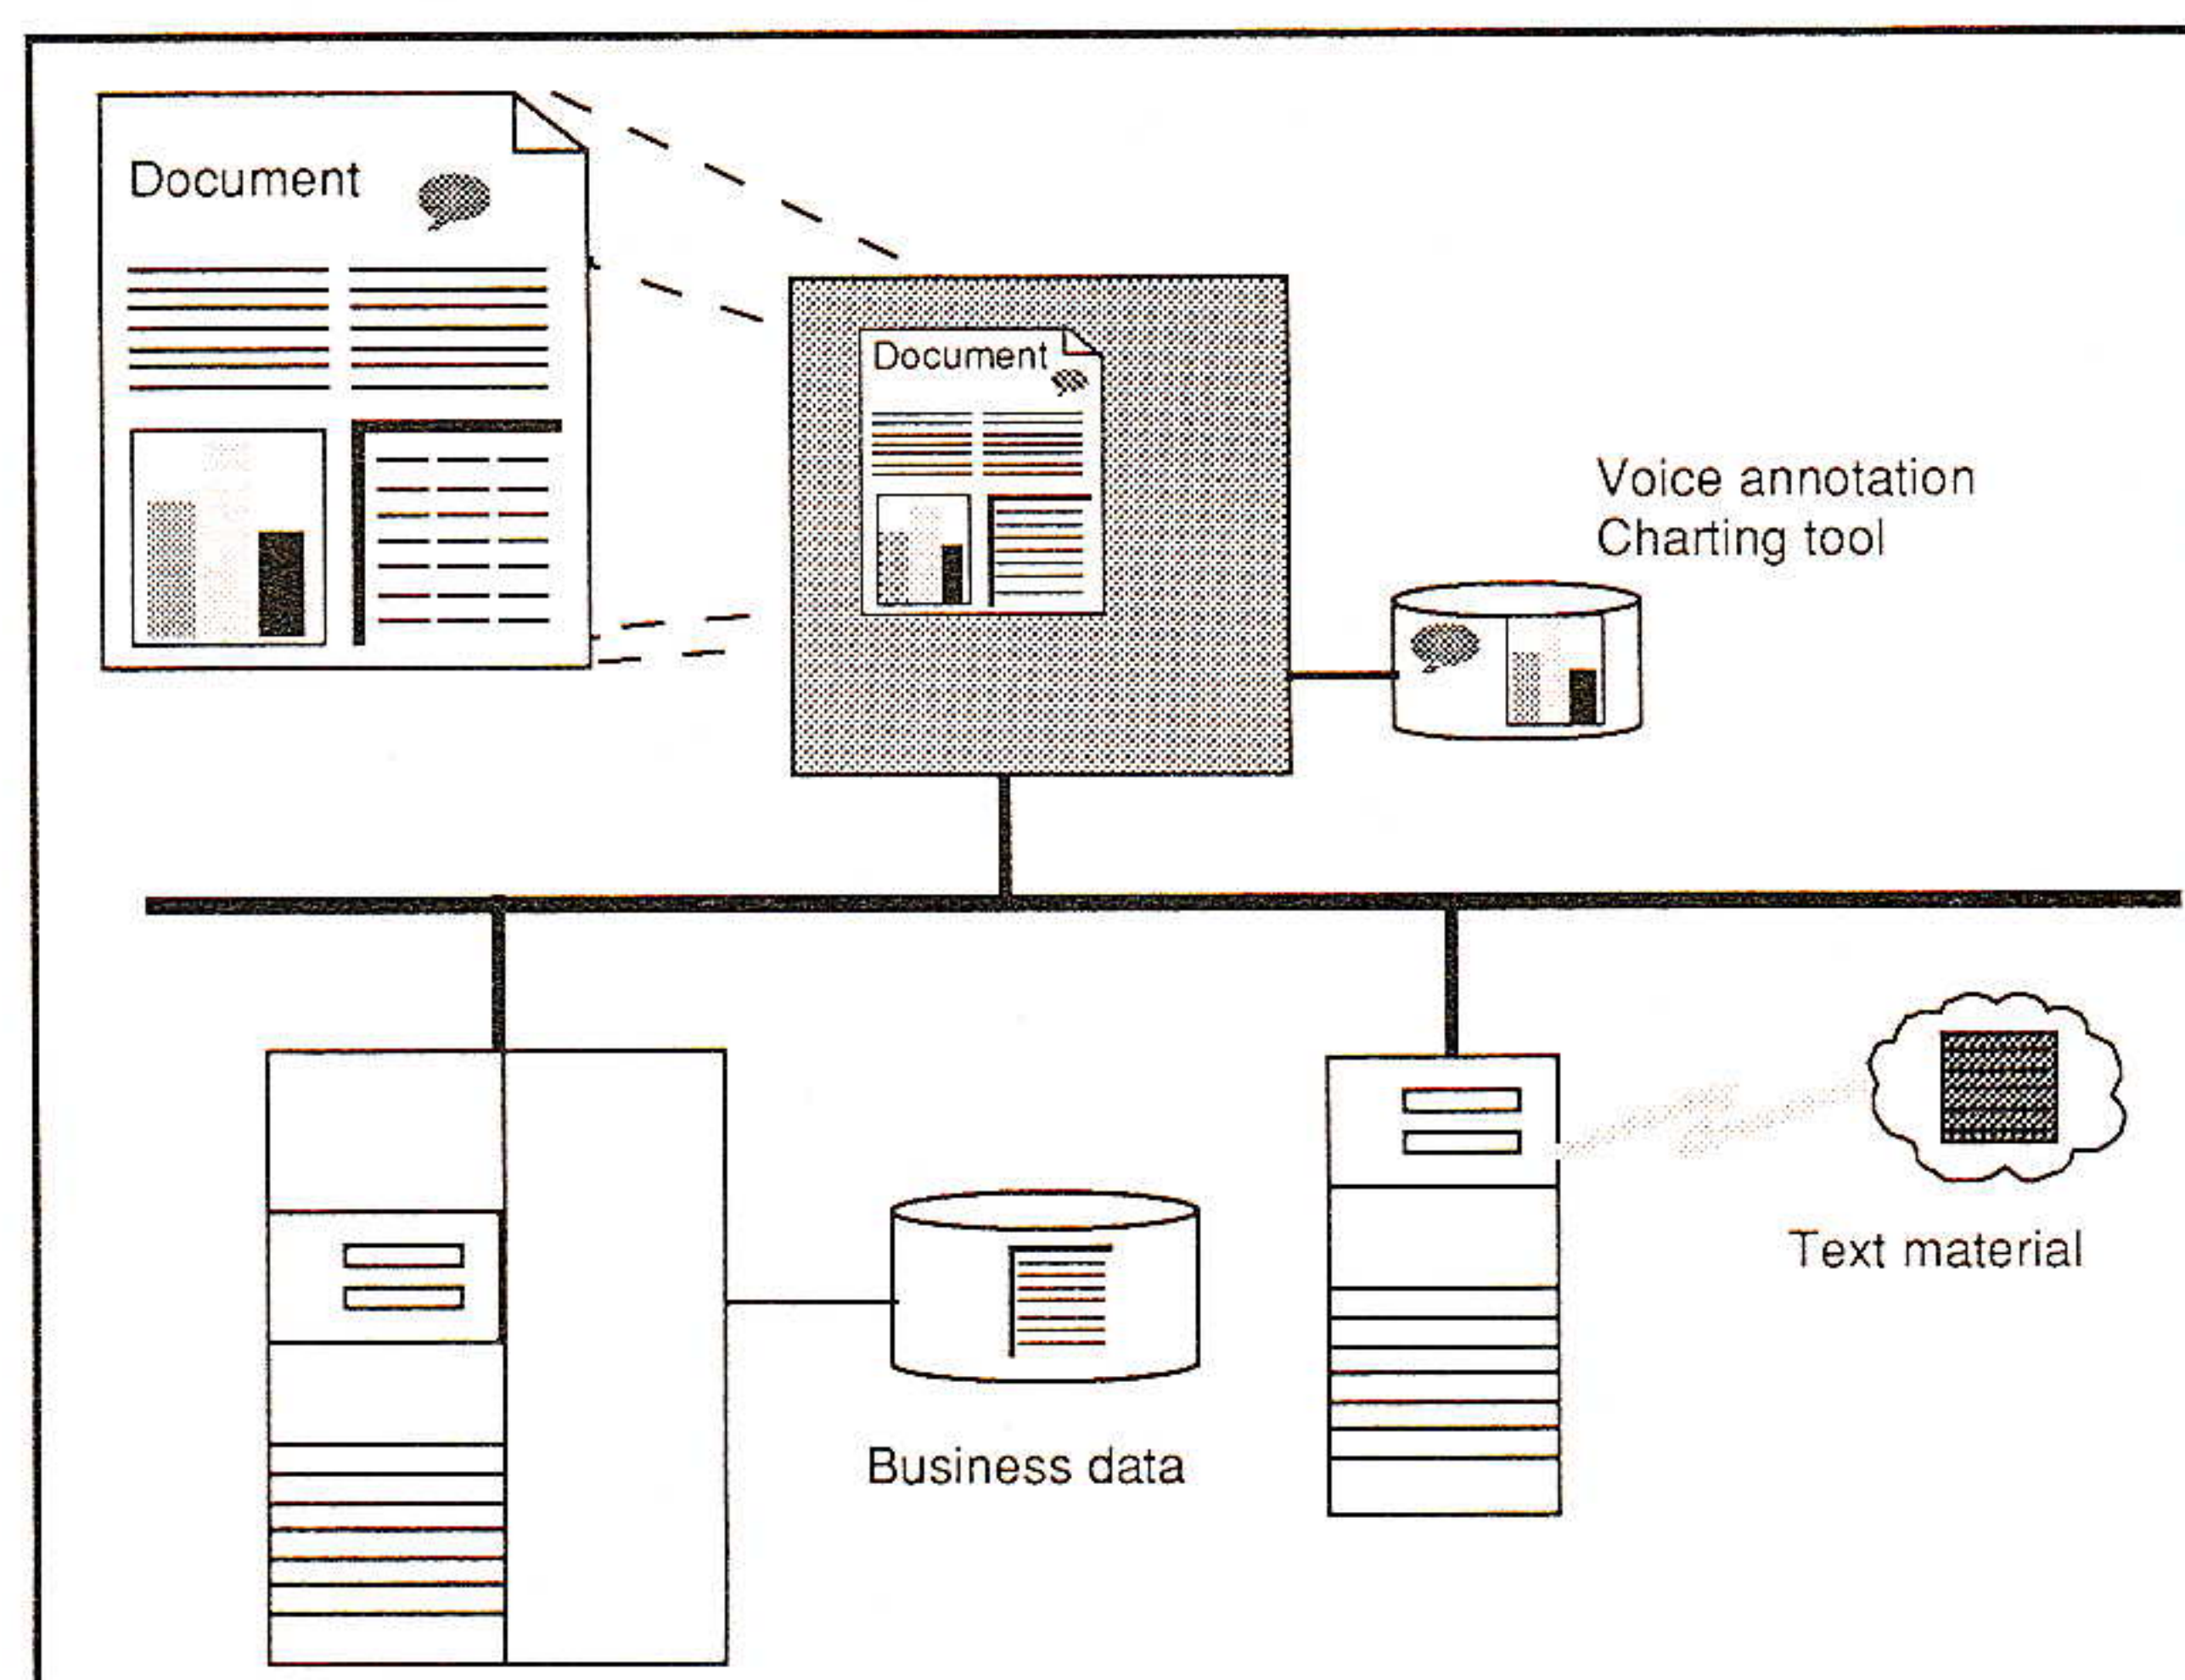
In an object-based environment, users can create compound documents consisting of multiple types of objects. A single "document" could contain text objects, spreadsheet objects, sound objects, project management objects and so on. This is essentially the document approach that Microsoft is pushing with its "Information at your Fingertips" slogan and its Object Linking and Embedding (OLE) technology.

But to be the most useful, such a compound document/object capability must be able to exist across a network. In a network environment, the compound document becomes the mechanism for accessing, integrating, analyzing, manipulating and then disseminating information and data throughout the enterprise. The elements of the document can continually update themselves, presenting the user with the latest information. Users could edit the embedded object if necessary. Other users working with an embedded instantiation of that object would then have those edits reflected in their work as well.

Ideally, such a document would integrate various objects from a variety of vendors based on a variety of platforms. The object environment eliminates the need to force users to learn a variety of interface styles and interactions for different applications. It also greatly simplifies the task of data integration. In short, it makes different types of data distributed across a variety of platforms accessible and useful.

Treating the compound document as a network service is the approach Digital is taking with its *Compound Document Architecture* (CDA), which is one of its Network Application Support services.

Such compound objects are good for more than just creating something for publishing. The object-based environments will prove extremely valuable in areas such as decision support. Consider an object-based visualization tool, similar to the new *Iris Explorer*. A user could integrate third-party modules into the basic tool, providing him or herself with the type of data visualization capability required for a particular task: plotting multivariate business data, for example. The data access itself would be handled through object-to-object communication.



The document consists of embedded objects located across the network as well as on the local workstation. In order to maintain the integrity of such a document, the links to remote objects must be persistent.

Figure 2: The networked document

CASE

In the object environment, source code could link to document specifications, test cases, branch coverage maps, performance analysis tools and the executable derived from the source. Having the executable object linked to the source code could allow the developer to dive into a browser/editor should a runtime error occur during development.

continued on next page

Networks and Objects (*continued*)

Digital is advancing such a platform with its *COHESION* software development environment that includes the use of the Application Control Architecture (ACA) Services. ACA Services are implemented as an object-oriented layer above operating system and network services, and provide facilities for handling application interaction, communication and start-up within VMS and Ultrix environments. Through ACA Services, Digital will encourage the integration of CASE products from different vendors.

Network Management

Combining the concept of network objects with the ability to create a compound object linked to network management tools, charting objects, relational or object-oriented databases and so on will be a powerful use of distributed object technology.

Current and historical network topology data can be kept in one tool and analyzed by another. A spreadsheet tool could extract network statistics from a network management application for analysis. Again, one of the major benefits delivered by the object technology is the ability to integrate easily a variety of tools from a variety of vendors.

Building links between existing applications

The path to object technology is one of migration rather than revolution. Hewlett-Packard recognized that when it started out with *NewWave* several years ago. *NewWave* delivers an object-oriented environment on top of Microsoft *Windows* through its *Object Management Facility* (OMF). The OMF binds applications with data to create large-grained objects (at the file level). The links between *NewWave* objects are persistent. The OMF tracks the movement of objects from place to place (on the local workstation) and maintains the link.

The *NewWave* environment offers direct manipulation of objects. Dropping a document icon on top of a printer icon causes the document to be printed, for example. Users of *NewWave* can create compound documents by linking or embedding objects of one type within objects of another. *NewWave Write* documents, for example, can contain voice objects, video objects, spreadsheet objects and so on. Embedded objects can update themselves. Users can mail *NewWave* objects and documents to other *NewWave* users. (These users would have to have the same applications or tools resident on their workstations to be able to utilize the objects.)

Once an application is written to the *NewWave* APIs and uses the necessary OMF calls, it gains instant integration with other *NewWave* compliant applications. However, recognizing that the industry wasn't about to shift to a new development platform overnight, HP gave *NewWave* the ability to support existing applications through *encapsulation*.

Encapsulation essentially builds a software shell around a non-*NewWave* DOS or *Windows* application to allow it to function in the *NewWave* environment. The shell provides communication between the application and the OMF. Encapsulation provides some of the following features:

- Binding a file to an application to create a manipulable *NewWave* object on the desktop.
- Cut and paste between DOS applications and other objects on the desktop.
- The ability to include older application in agent tasks.

The NewWave *Agent* facility is essentially a software robot that provides cross-application procedural automation within the environment.

There are limitations to NewWave that raise the hackles on some object purists. For example, there is no inheritance under NewWave, partly because the "objects" are so large (applications and files). However, it was and remains a very pragmatic approach to providing a migration path to object technology. HP also plans (along with its development partners) to move to a finer-grained object model with full support for all the attributes of a classic object environment.

Currently, NewWave is limited to the DOS/Windows platform, and does not support objects distributed across a network. NewWave does offer the ability to store objects on servers. With the advent of UNIX and OS/2 NewWave clients, HP will be greatly enhancing the capability of the client as well as NewWave's ability to function in a distributed environment. (More on the emerging Distributed Object Management Facility later.)

HP had hoped that NewWave would become the de facto standard on the desktop for the object environment. HP has licensed NewWave to other vendors (such as NCR, DG and AT&T) and helped found the *Object Management Group* (OMG) to push for standardization on aspects of object technology. HP really was the first vendor to insist that the industry needed a common platform of object technology in order to move ahead.

Microsoft has been a major barrier to NewWave. Although the environment nicely augmented Windows, and although Microsoft agreed to support NewWave with *Excel*, Microsoft had its own technology plans that precluded buying into NewWave wholeheartedly.

Microsoft and OLE

Microsoft's new *Object Linking and Embedding* (OLE) protocols represent Microsoft's initial step at providing an interoperable, object-oriented environment similar to NewWave on the user's desktop. [Ed.: They didn't consult with me before choosing this name!]

OLE defines some extensions to the Windows *Dynamic Data Exchange* (DDE) protocols and the clipboard that allows users to create compound documents out of various application objects. You could embed a graphic in a text document, for example, and retain the ability to edit the graph simply by clicking on it within the text document. Or, you could have the graph itself linked to a set of numbers in a spreadsheet.

OLE does not deliver full in-place object editing. Clicking on an embedded object launches the appropriate application, which then brings the target data up in a window.

Microsoft provided an earlier, less functional foundation for such capabilities with the original DDE protocol in Windows. Indeed, NewWave provided more functionality several years back (OLE provides 2 percent of NewWave 2 years later, one critic sneered.) But Microsoft is not re-inventing NewWave by implementing a different type of Object Management Facility. Microsoft is simply taking a different approach (more on this below).

Networks and Objects (*continued*)

You may or may not agree with Microsoft's approach in charting its own course independent of the OMG and NewWave. However, there are several attributes of OLE to keep in mind:

- OLE will be delivered on Windows, PM and Macintosh. Every graphical application Microsoft delivers from now on will support OLE. Notably absent is any discussion of implementing OLE libraries on UNIX.
- Microsoft has paid attention to other voices in this. OLE is the synthesis of proposals and specifications generated by folks at Microsoft, Lotus, Aldus and WordPerfect. Those four vendors represent quite an impressive share of the desktop application market, and those four plan to implement OLE. Lotus is using it in *Notes*, for example. This provides the Microsoft OLE technology with something NewWave has yet to achieve—a suite of leading applications exploiting the technology.

With these forces behind it, OLE is destined to be a de facto standard. Software developers better pay attention to it, and systems vendors better figure out how to support it or encapsulate it in their own strategic object-oriented environments. Sensibly, HP plans to support OLE within NewWave.

OOFS

OLE has weaknesses, particularly in distributed object management. Essentially, there is no good persistent mechanism for tracking a link to a remote object. Should someone else move the remote object you've referenced, for example, you have no way of knowing where it went or what it became. Microsoft recognizes this limitation, and promises to address it with the development and release of its *Object Oriented File System* (OOFS). Even in the most optimistic of Microsoft's scenarios, this reliance upon OOFS pushes distributed object support out a minimum of two years. For some industry veterans, this smacks of hand-waving and the smoke and mirrors of Windows 1.0.

What Microsoft offers with OLE, while a localized solution, still is quite good. The prospect of having a suite of major applications supporting it is even better. And the promise of Macintosh and PM support is even better yet.

Being able to link and embed objects across Mac, OS/2 and Windows platforms will be a definite boon to users. In the short term, however, such cross-platform integration will prove a bit clumsy, as there is no distributed object management solution.

We should also point out that OLE is concerned with application integration, not with compound document interchange across heterogeneous systems. OLE is a application environment solution, not a document interchange architecture. To be sure, you can take OLE documents across systems, but you will need the appropriate applications on the other end.

Model

OLE is designed to allow users to create and edit compound documents using multiple kinds of applications. Users can move objects from one place to another, and may link those objects to other locations. Objects have content and behavior.

The OLE semantic model has two basic elements: *objects* and *containers*. The elements have two fundamental relationships—containment and reference—and support three basic operations: move, copy and link. Containment does not affect the content or the behavior of an object.

Many different types of things can be OLE objects: charts, drawings, tables, videos, voice annotation and so on. Any type of application document can contain or link to any other object. As is NewWave, the OLE environment is open. Theoretically, a vendor can implement a voice annotation application, for example, which all other OLE applications then can use.

There is a fundamental distinction between linking and embedding. A *link* provides a dynamic reference to another object. Designed for sharing information, linking provides a two-way dynamic update capability with the original object. Should you link to a graph from your document, for example, a change in the graph will appear in your document. Similarly, editing the graph from within your document will change the original object, and those changes will be reflected in whatever other links the object supports.

Embedding essentially clones an existing object. Although the object is still editable, any changes to it will not appear in the original object. Lotus seized upon this capability as a way around the lack of persistent linking.

In *Notes 2.0*, Lotus embeds objects in documents, and then replicates the documents across other Notes servers. This becomes a quite functional way of distributing compound information.

Implementation

Microsoft built OLE on top of the current DDE and clipboard mechanism. With those as underpinnings, the user interface conventions for building OLE-based compound documents are based on the cut, paste and paste-link commands familiar to many Microsoft application users. This is quite different from the NewWave environment that allows users to pick up an icon representing one object and to drop it on another to embed it. As Microsoft implements the drag and drop convention in Windows, it will implement it in OLE, too.

For the developer, especially for developers who have already worked with DDE, implementing OLE is not a major burden. Microsoft provides a linking and embedding (L&E) subroutine library further to ease the task. Presentation data is cached transparently for fast display of the components.

OLE uses a client-server model. The client is the container document and provides storage space for the object, and place to display and print, and user commands to open the object for editing. The server is the editor of origin of the contained object. It supports the object and provides a persistent representation, a presentation, and editing tools.

The client uses a *function library* to manage the server object. The client calls the library to load or save the object, to display or print the object, and to open the object for editing. In turn, the client supplies the library with functions to read and write data, access to its display space and callback functions to notify the client of changes to the object.

The client library can launch the server if necessary. The server also uses a function library to manage client communication. The server registers itself with the library, provides function to create documents and to set and retrieve data and renderings, and notifies the library of changes to the object.

Future directions

OLE is only one of Microsoft's technologies for providing compound documents/objects. Microsoft also has three other major pieces in the works.

continued on next page

Networks and Objects (*continued*)

- *Object-Oriented File System*: Microsoft intends these future extensions to the file system to provide link maintenance and change notification (locally and across the network). Microsoft, in other words, intends to embed its object management capability within the file system itself, rather than create an external construct such as NewWave's Object Management Facility. This is currently a major design difference between the two approaches.

Any OOFs object will be persistent and shareable. Microsoft is seeking to provide a uniform way to find and enumerate all kinds of objects, with standard protocols for load, save, display, edit and print. They are hoping to be able to allow a user to search and find an object within a document via this mechanism.

- *Object-Oriented Shell*: Window's future (Windows 4.0) is to become an object-oriented shell along the Xerox *GlobalView* or *Metaphor* lines.

- *External Macro Programming Language*: Microsoft knows that it needs a cross-application macro programming language. OLE objects provide an interface for access to data and subsequent manipulation. A macro language would take this further with visual programming tools and a visual macro recorder.

Digital ACA

Both NewWave and OLE suffer from the inability to provide distributed object management. HP, in conjunction with Sun, is in the process of developing a Distributed Object Management Facility to provide a foundation for such distribution.

But Digital already has a facility for linking application-objects across a network: the *Application Control Architecture* (ACA). ACA and ACA Services are the foundation for a variety of existing capabilities:

- The *LiveLinks* between object types in Digital compound documents
- The *Builder* facility in *DECdecision*
- CASE tool integration with *COHESION*

Digital, recognizing an opportunity, is integrating OLE and DDE with ACA. The result of this venture will allow Windows OLE and DDE applications to exchange data with VMS and UNIX applications on Digital systems. This in effect adds the missing distributed object management capabilities to OLE.

ACA defines an object-oriented approach to application integration in a distributed environment. Although it does use an object-oriented model, ACA does not require supported applications to be written in an object-oriented language.

There are three levels of application integration with ACA Services. At the first level, ACA Services registers and invokes applications anywhere on the network with no modification to the application. At the second level, ACA Services can publish all or a subset of the externalized functions of applications with existing command line or call interfaces. Again, this requires no modification to the application code. The encapsulating message resides outside the application. Once a program has been defined to ACA Services in this manner, any other application on the network (with authorization) can use the externally registered application functions. At the third level, applications can incorporate calls to ACA Services.

This allows applications to conduct ongoing dialogs with other applications, or to function as application servers. ACA thinks of applications as groups of services, or methods. Classes represent the external behavior of an applications. There is a set of methods associated with each class that operates on instances of that class. ACA currently defaults to file-level objects (as in a typical compound document or DECdecision application) but the framework is there for greater granularity.

Once a class is defined as a parent class, ACA supports inheritance from those classes as well as dynamic binding. ACA specifies services for:

- Application definition
- Application naming
- Dynamic application location and invocation
- Binding of information objects to applications
- Communication between applications

ACA Services thus is providing a model for application interaction and management, rather than managing objects themselves. In other words, ACA Services can call applications' methods, can bind dynamically or support inheritance, but it doesn't know anything about what those methods are acting upon.

With its support for both existing and new applications, ACA provides a nice platform to bridge into the next area of discussion: developing new object-oriented applications.

Writing for the Distributed Object Environment

There isn't a lot of debate over the essential functions or mechanisms of an object-based distributed system. Drawing from a variety of academic and research projects over the years, we can summarize these services as the following:

- Creation and deletion
- Naming
- Location
- Remote Method Invocation
- Security
- Control of concurrent execution
- Serializability
- Failure atomicity
- Control recovery
- Exception handling
- State management for short-term and long-term storage

Some of these elements can—and should—be handled by the underlying network environment. However, as we noted above, there is a need for some form of distributed object management facility to support the bulk of these service elements.

Here we will discuss three approaches to solving these problems: The joint HP/Sun *Distributed Object Management Facility*, Sun's interim *ToolTalk* solution, and NCR's *Cooperation*.

continued on next page

Networks and Objects (*continued*)

To this list, you should add Digital's ACA (discussed above) as well as technology from a variety of smaller companies and consortia. In the initial rounds of submissions to the OMG for an *Object Request Broker* technology, some of the large vendors mentioned above partnered with smaller vendors. Digital, for example, united with HyperDesk (a spin off of Data General's object program purchased by ASCII).

DOMF

Surprising many in the industry, HP and Sun joined together to create a *Distributed Object Management Facility* (DOMF) to support distributed object-based systems. Previously, the companies had been at competitive loggerheads about the mechanisms used in their respective distributed computing environments: NCS (and now the OSF DCE) for HP; ONC for Sun.

The new DOMF announcement marks an end to squabbling over RPC mechanisms with a recognition that customers want to have a consistent application interface to a variety of network mechanisms.

We must point out that Digital's ACA, as well as other technology submissions to the OMG, *also* are network independent. But the partnering of Sun and HP constitutes a very public recognition that DCE platform independence is important for the future.

Emphasizing the platform-independence of the DOMF, Sun and HP each will ship an RPC-specific implementation for its own RPCs, along with an RPC compiler for either NCS (HP) or ONC (Sun). Further keeping in line with each vendor's choice of a distributed computing platform, the two systems use different naming schemes (unique object IDs for HP vs. names for Sun), and in their optimization for object size and granularity. Sun favors very small objects, while HP favors large. Each also uses separate security mechanisms: HP uses DCE's security and RPCs; Sun uses ONC's *Secure RPC* and will later be able to use *Kerberos*.

However, HP and Sun do share a common vision of the DOMF consisting of a manager of object managers, making it easy to accommodate different local object systems.

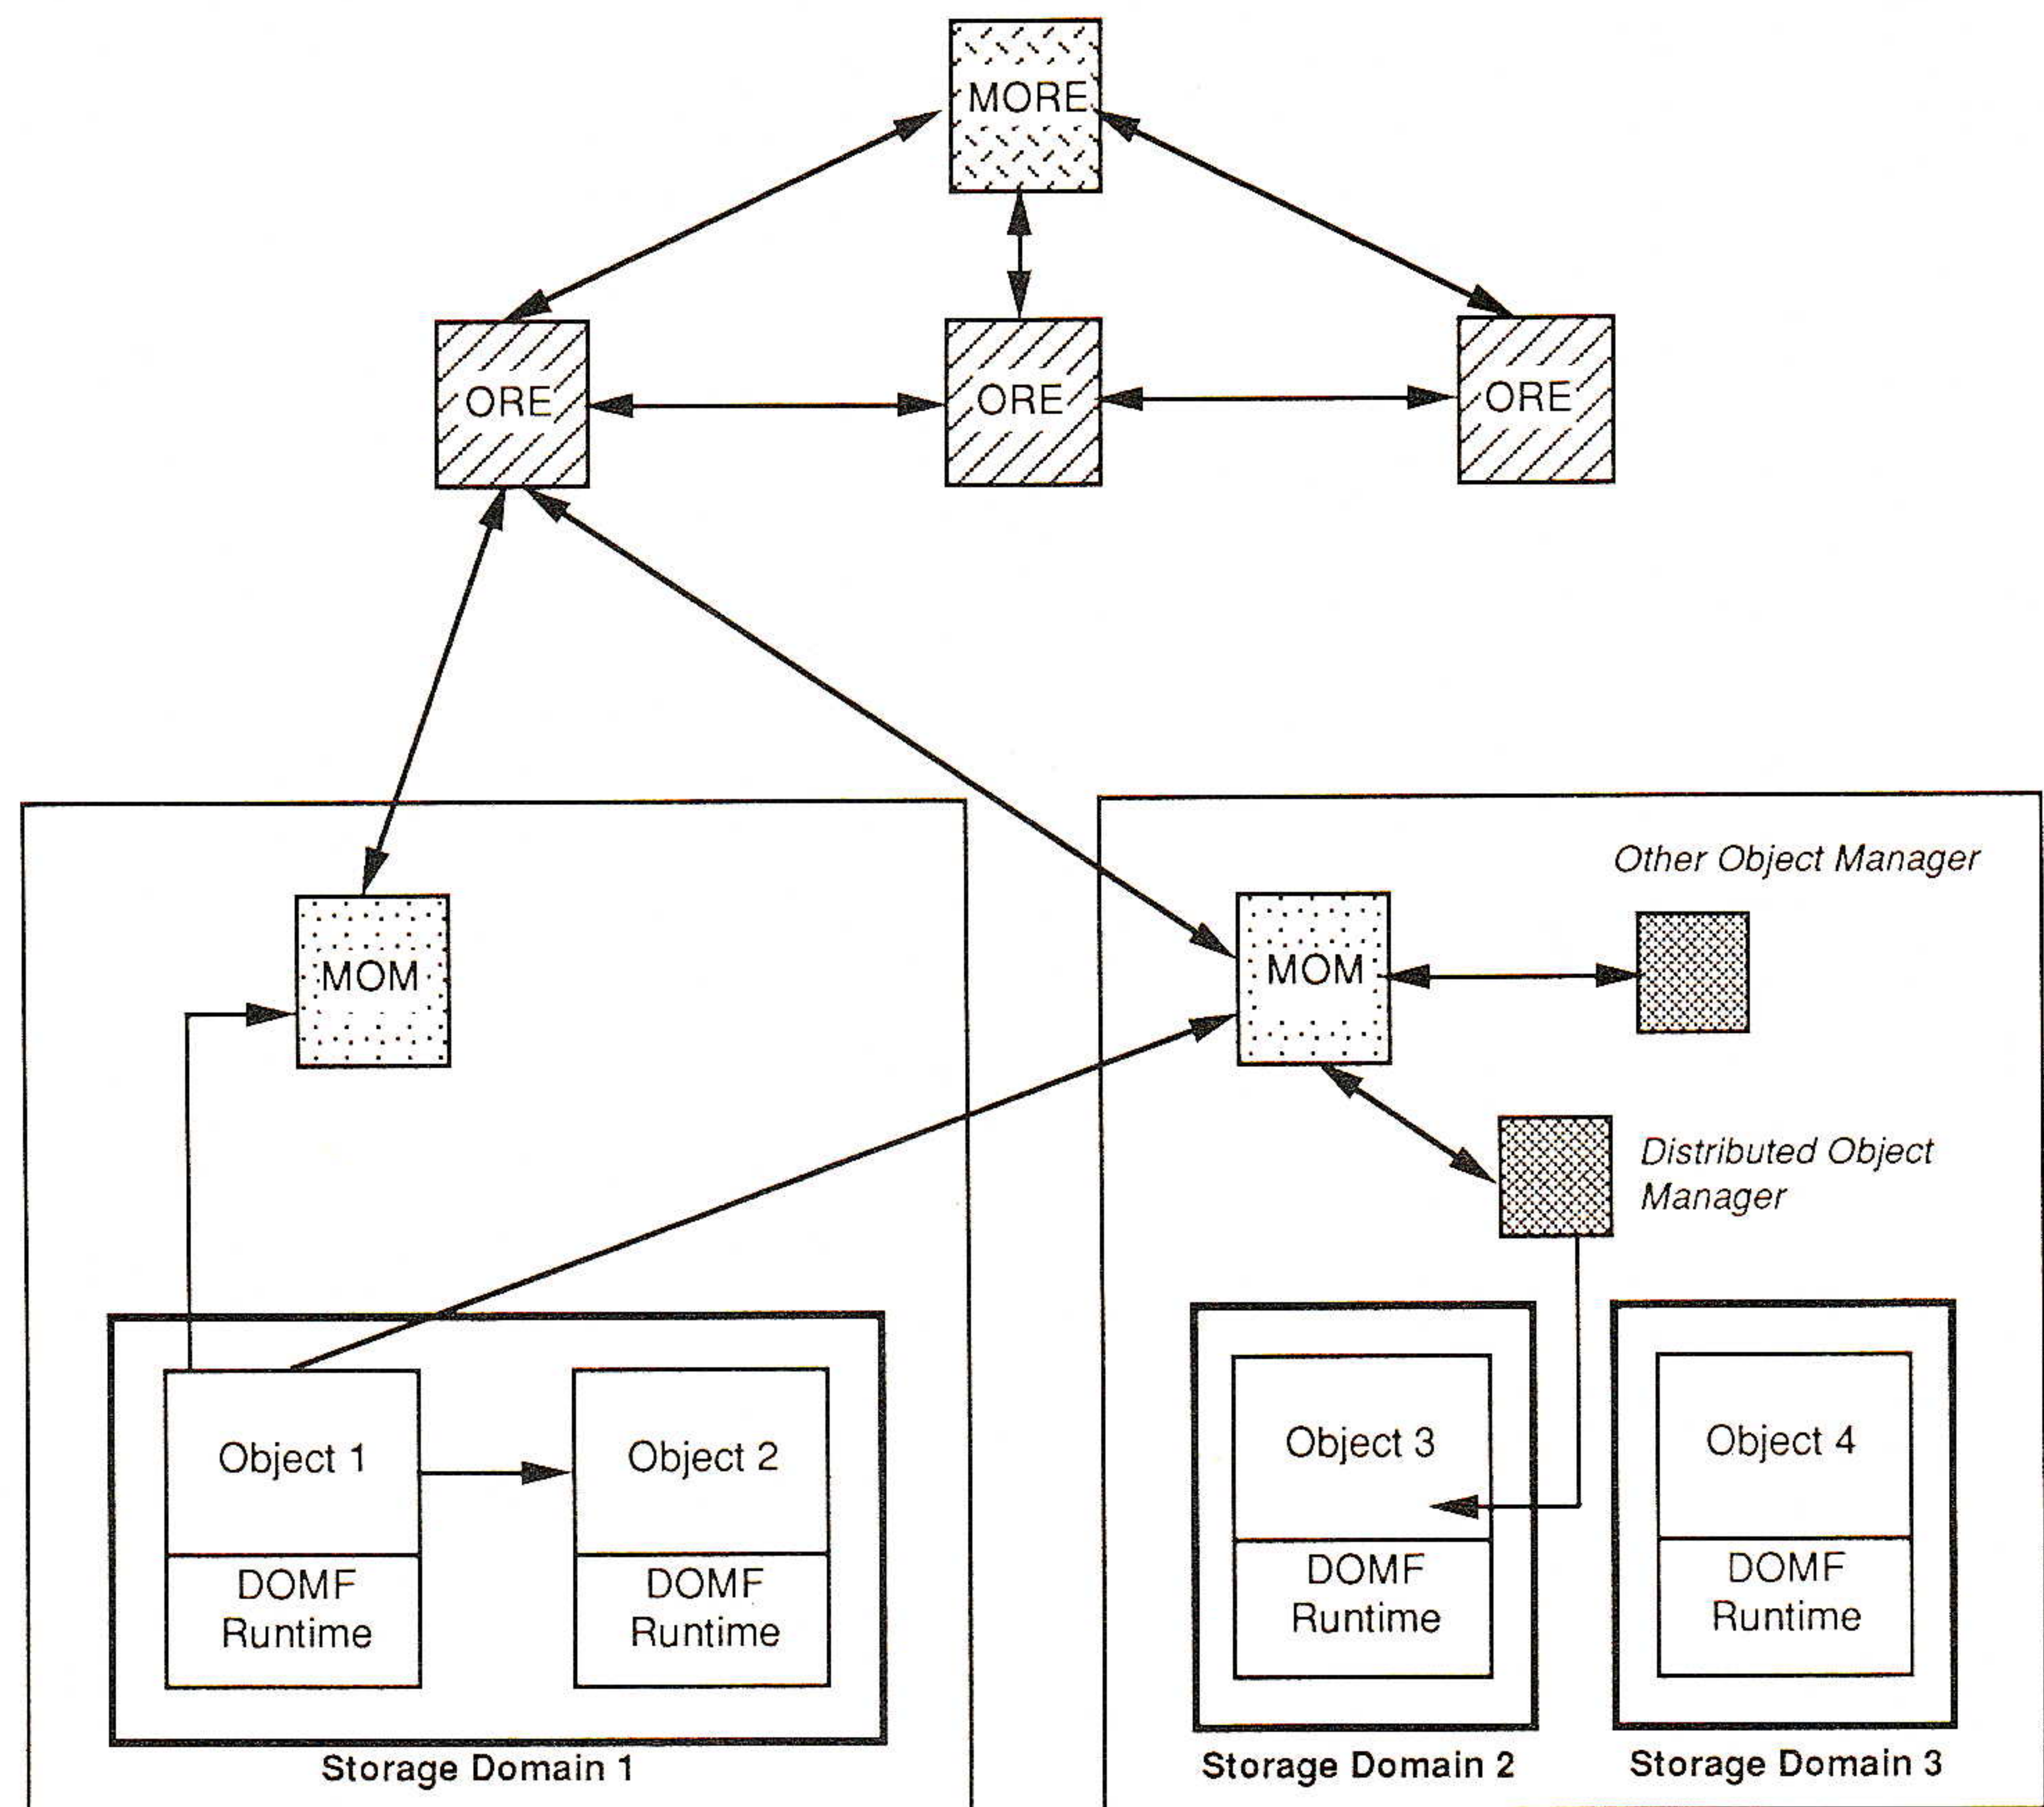
The DOMF is modular, and uses defined interfaces to services such as X.500 or a local object manager. Registering objects with the DOMF requires the use of a *Class Definition Language* (CDL), or of an API. CDL is the first step toward a common, single interface across both HP and Sun environments. Registration is a formal process, and more static than ACA Services.

ToolTalk

Sun also recently unveiled *ToolTalk*, a messaging protocol and message delivery service that provides a transitional step toward the DOMF.

The ToolTalk toolkit provides both multicast and object-oriented messaging between UNIX applications written by different vendors. In this light, it performs a function comparable to that of DDE/OLE in Windows or ACA Services from Digital.

ToolTalk is responsible for delivering messages to the intended targets. It can also launch an application to satisfy a request for an object if required. To use ToolTalk, developers must incorporate calls to the ToolTalk API.



The DOMF defines a hierarchical structure for the management of object regions. In the example above, Object 1 is sending a message. If the message is for Object 2, within Object 1's storage region, the message goes direct. To send a message to another object (Object 3) in a different storage region, Object 1 resorts to a Manager of Object Managers (MOM). If Object 1 has sent the message before, it may go directly to the MOM in charge of the recipient's storage domain. In the other case, Object 1 sends the message to its MOM, which it turn passes the request up the hierarchy to Object Region Experts (ORE) and MetaObject Region Experts (MORE). There are also Sibling MOREs (SMORES) at the top level. One of these elements will find the appropriate MOM, and pass the message to it. That MOM will in turn pass the message to the appropriate object manager. It may seem convoluted, but the architecture is designed to support millions of nodes.

Figure 3: The HP/SUN DOMF.

The ToolTalk API runs on each machine in conjunction with a ToolTalk server. The server functions as a daemon that sits and listens for messages and determines the targets for which they are intended. Developers register message patterns with the ToolTalk Service. The service scopes out the pattern of incoming messages and matches them to the registered receivers.

This year, the *CAD Framework Initiative* will be using ToolTalk as its messaging facility in this year's Design Automation Conference demonstration of inter-tool interfaces.

Networks and Objects (*continued*)

NCR NCR's distributed object management technology is part of *Cooperation*, its applications delivery and development environment. NCR based Cooperation on NewWave clients, but on its own object model and distributed computing foundation.

NCR took a harder-nosed approach to its object manager than did others (such as HP). It handles other object managers, but mostly as encapsulated foreign bodies.

NCR defines *all* resources on the network as objects, and gives each Cooperation object its own methods for rendering its messages and arguments into streams for communications and storage. NCR does not use an RPC mechanism, but instead uses a *Remote Method Invocation* protocol, which is dynamically linked to the object itself rather than managed as an external scheme. This is a major difference between NCR's approach and that of Digital, HP, Sun etc, which base their scheme on the use of their RPCs. Currently, NCR is using IBM's APPC, X.25, and LAN Manager as transport protocols. Cooperation requires only the basic data transport services for implementing RMI-based communications in the distributed environment.

Cooperation provides a *Procedural Programming Interface* (PPI) that allows the use of Object Oriented languages other than C++ via customized bindings. The PPI also allows procedural languages to exploit the full object oriented mechanisms of the Cooperation environment.

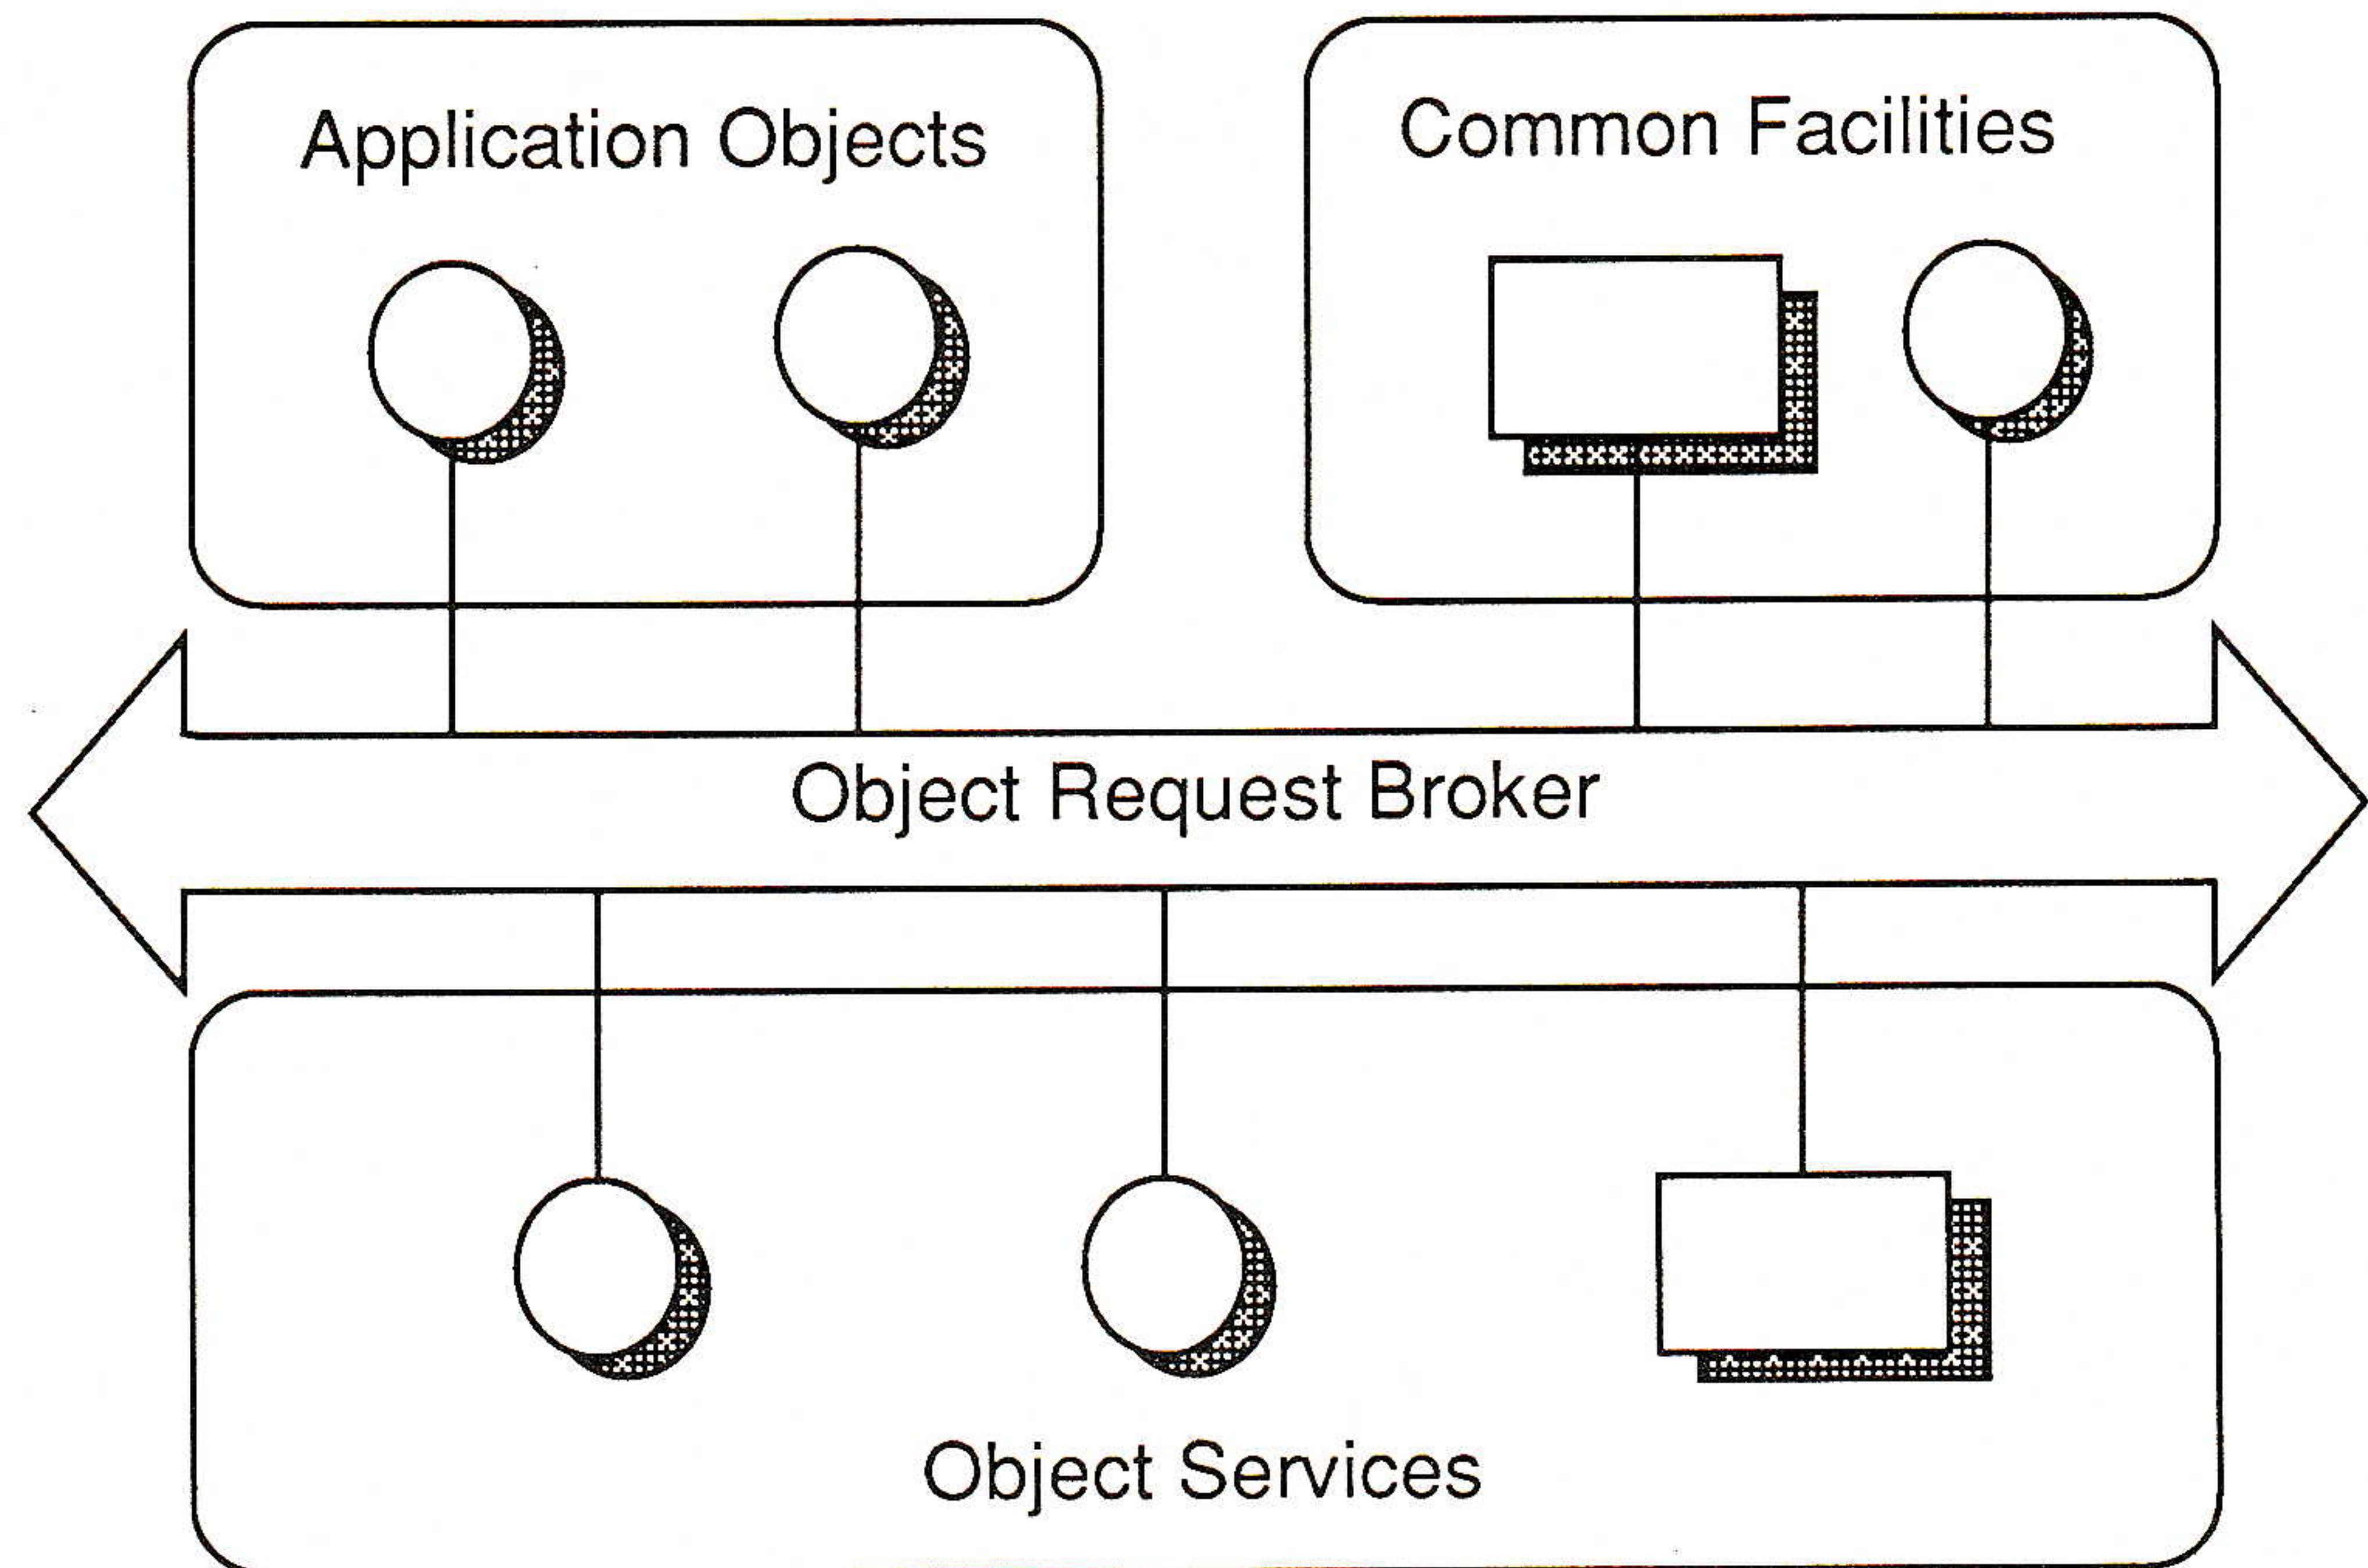
Towards standard interface definitions for objects

The industry faces a dilemma. It is way too early to expect any sort of widespread consensus on a common object model. Even the vendors with the most experience in the area can't agree. Yet without some form of standardization, we face a technological Balkanization. That would run counter to the mandate for open, interconnected systems. The solution appears to be to agree up front on a set of interfaces that would allow different objects and object managers to communicate and interoperate. The definition of that common set of interfaces is the task the Object Management Group has assumed for itself.

Unlike the OSF, the OMG doesn't build or sell software. Instead, it spends its energies in hammering out a set of interfaces for managing objects in the distributed environment. In this effort, it relies almost completely upon the efforts of member companies. OMG per se has no technical staff, aside from a very astute technical director who goads the members of the technical committee along the proper paths.

OMA The OMG published an *Object Management Architecture* (OMA) in October 1990. The OMA contains the foundations of the future hoped-for standards, including an abstract object model and a reference architecture. The OMA consists of four main facilities:

- The Application Objects
- Common Facilities
- The Object Request Broker
- Object Services



The OMA consists of four elements: Application Objects; Common Facilities, the Object Request Broker, and Object Services.

Figure 4. The OMG Object Management Architecture.

Application Objects are the non-standardized applications that will use the standardized facilities of the other three elements. Common Facilities represent commonly used application systems and standard interfaces to common data, such as:

- Mail systems
- Spell checkers
- Control system interfaces
- Input/output interfaces
- Agent facilities

These are software elements that once would have been a discrete application, but have become used as horizontal tools for the entire system. Application Objects will become common facilities once they achieve the status of de facto standards.

ORB The *Object Request Broker* (ORB) is the core of OMA. It is responsible for managing the communication between objects, including location, naming, delivery, activation, synchronization, exception handling and security.

OMG mandates that the ORB be network independent, and that it not be involved directly with the structure, form or capabilities of the objects on the network.

One of the ORB's main tasks is keeping track of the objects (a role similar to that of the OMF in NewWave). The ORB does not specify a name service; it will use the name service capability of the underlying platform if necessary. But the ORB must register objects. Registering objects allows the ORB to keep track of the available services and interfaces.

continued on next page

Networks and Objects (*continued*)

The OMG has issued a request for technology for the ORB, to which a number of vendors, including DEC, HyperDesk, HP/Sun, and NCR and have responded. The various vendors chose different mechanisms for registering objects:

- Explicit registration of remote objects using a class definition language. This creates a somewhat static environment for distributed object management, in that adding a new object class or changing an existing class requires the redefinition of existing classes and the recompilation of all stubs.
- Registering remote objects via an API to the broker. The API-based approach allows changes without recompilation, and hence supports a more dynamic environment.
- Automatic creation of the class hierarchy through the use of a special preprocessor when the objects are compiled.

Early in the summer, the technology selection for the ORB had narrowed to two camps: an HP/Sun/NCR grouping pushing the explicit registration method, and a Digital/HyperDesk group advocating the API-based approach. Just prior to the OMG's June 3rd deadline, both groups decided to work together to create a single unified specification. At the time of this writing, OMG extended the deadline until August 20th. Should the submitters fail in their efforts, the OMG will launch a new selection process.

The final element, the Object Services, provides the basic functionality of object management: object structure, creation and deletion, transaction control and security protocols.

Final comments

Conceptually, the melding of objects with distributed computing makes a great deal of sense. It promotes a location-independent view of data and services, isolates applications from underlying mechanisms and protocol subsystems, and promotes flexible application architectures. It also *should* make it easier to develop solutions for the distributed environment.

When such an environment fully arrives, it will have caused a number of changes in the industry. The advent of a true object-based environment will open up new markets for class libraries, or for object tools that will augment other applications. Such an environment would encourage the incorporation of a particular visualization module, for example, as a supplement to a regular decision support package and so on.

There is tremendous ferment in the industry right now over the shape and nature of the enabling technologies that will create this environment. The OMG has almost 200 members and interested parties, including all the major systems vendors as well as large user organizations. But the OMG appears to be headed toward defining a way for multiple object models and managers to work together.

The vendors responding to the OMG RFT have been clear that they intend to press on with their own technologies, incorporating the ORB interfaces as possible and appropriate. Each has said that they could accommodate a technology selection that was not completely their own.

Other members of the OMG will push ahead with their technology projects as well. Microsoft is of course a key player, perhaps even more so now that it has acquired the services of Paul Leach, formerly of HP/Apollo, a founder of Apollo, and designer of NCS. IBM continues to forge its own path ahead, although its presence at OMG meetings means that it takes the standardization effort seriously.

IBM seems determined to proceed on its own course in developing a platform-independent, object-based distributed environment. It first created the Patriot Partnership with Metaphor with that goal in mind. It then entered into a strategic technology arrangement with Apple, one aspect of which should be the creation of another joint venture to pursue such advanced object-oriented work. IBM then acquired Metaphor and threw the Patriot technology into the Apple-IBM pot. This could be an extremely significant platform, and yet it is an unknown one at this point.

Although the nature of the technology enabling the bulk of object oriented distributed systems might not be clear, the projected effect on the market is a little more so. In an environment supporting modularity, integration and reuse, the dynamics of the software market are going to change radically. Vendors who used to charge a lot per unit for a few units of software are going to have to mentally switch to charging less per unit for a much greater volume. Dave Liddle, Metaphor's president and the head of Patriot Partners, always claimed that this shift in dynamics would greatly benefit the smaller software vendors, who now would have the opportunity to sell their expertise as integratable modules.

The transition to these environments may take the rest of the decade, but the time to start paying attention to the issues is now.

References

- [1] Microsoft, "Object Linking and Embedding—Extensible Application Protocols (pre-release version)," December 1990.
- [2] M. Aksit, J. W. Dijkstra, & A. R. Tripathi, "Atomic Delegation: Object-oriented Transactions," *IEEE Software*, March 1991.
- [3] Digital Equipment Corporation, "Application Control Architecture," 1991.
- [4] Digital Equipment Corporation, "Network Application Support Overview," November 1990.
- [5] Object Management Group documents: 91.3.4–91.3.17 (Overviews of submissions in response to the OMG ORB RFT).
- [6] Millikin, M., "Distributed Computing Environments: Laying the foundation for the future of interoperability," *ConneXions*, Volume 4, No. 10, October 1990.

MICHAEL D. MILLIKIN is a principal of Gunstock Hill Associates (GHA), the computer market research and publishing division of Technomic Publishing Company. GHA closely tracks the emergence of the new distributed computing environments based on open systems and networking. Mr. Millikin is the editor of *Business Visualization*, a monthly report on hardware, applications and network platforms for the visual analysis of business information. He will be launching the *Client-Server Report* late in 1991. Mr. Millikin also provides consulting on technology and market issues to both the vendor and user communities. In addition to his consulting and writing duties, Mr. Millikin has spoken widely in the US, Europe and Asia on technology trends, open systems, and advanced network environments. Mr. Millikin spent nearly a decade between Seybold Publications and Patricia Seybold's Office Computing Group, at which he was the chief technologist and a vice president. Mr. Millikin was editor-in-chief of both Patricia Seybold's *Network Monitor*, and an editor of both Patricia Seybold's *Office Computing Report* and Patricia Seybold's *Unix in the Office*. At Seybold, Mr. Millikin also wrote several lengthy reports on distributed computing technologies for the large corporate Enterprise.

At INTEROP 91 Fall don't miss session S27 "Object-oriented Approaches to Building Network Applications" on Thursday, October 10 at 3:30pm.

Multivendor Enterprise Networking A Case Study

by Wayne H. Russell, American Management Systems

Introduction

In recent years, the concept of *enterprise networking* has become very popular in the industry. However, when reality hits an organization, the political, practical and technological implications conspire to make the actual implementation of such a network an elusive goal. The case study presented here is a snapshot of where a large government agency is today in making multivendor enterprise networking a real solution to actual management and user needs.

Current environment

This government agency desires to install a data communications network to serve the Washington, D.C. headquarters facilities comprised of three large office buildings with approximately 15,000 potential users. As with many large corporations, there is a mainframe environment which houses many database applications. In addition, there is a "potpourri" of minicomputers located in various departments providing support for a wide range of applications from office automation to CAD/CAM. In recent years, the proliferation of PC LANs (approximately 10 vendor architectures exist) has increased with very little emphasis on interoperability. Various backbone network approaches to connect the multiple workgroup LANs are used in the three buildings, but these approaches do not provide universal connectivity or access. The level of network management varies from workgroup LAN to workgroup LAN from rudimentary server backups to more sophisticated monitoring of network traffic and diagnostic testing. However, the management of workgroup LANs is localized and there are no global standards or procedures.

To address the existing situation and to begin the process of transitioning to an enterprise network spanning the entire organization, the agency initiated a set of steps including requirements analysis, alternatives analysis and prototype development to show that the goal of multivendor enterprise networking could be achieved.

Identified problems

An initial analysis of the organization revealed the following basic problems with the existing networking mechanisms:

- No standard wiring plan or mechanism for intra- or inter-building connectivity
- No common standards for networking protocols
- No standards for network management
- No standards for electronic mail or file transfer
- No ability to access applications residing on a remote network server (e.g., Novell) that differs from the local network server (e.g., Banyan)
- No standard mechanism to access local and remote mainframes
- No means to allow a user to move from one location to another and still have access to his or her "home" server
- No plan for implementing all or parts of the *Government Open System Interconnection Profile* (GOSIP)
- No high speed backbone network to support planned CAD and image applications and file transfers

In addition, two issues of paramount importance to the network users and agency management are: (1) becoming a network citizen without giving up the established hardware and software base (i.e., the user's concern) and (2) the establishment of a target network architecture to which user's could evolve to over time (i.e., management's concern).

Target architecture

The resolution of the previous two issues and a solution to the identified problems was the basic task. The key word in building the multivendor enterprise network is "interoperability." The following is a quote from the editor of LAN Magazine contained in the Fall 1990 issue entitled *A Guide to Interoperability*.

"Interoperability is the substrate upon which applications will run that can change the way corporations do business. Interoperability means fluent conversations. Users want to access corporate databases that may span several continents. Users want to exchange information and messages with their fellow employees, customers and suppliers, regardless of where they are located. The need for distributed, mission-critical applications is real. They need to be built as well as deployed.

The road to interoperable computers and networks is rocky. But the edge a mission-critical application can give a company can mean the difference between success and mediocrity."

The previous quote essentially defined our task. To come up with the "target architecture" envisioned by the agency management, we divided the network into several design areas and examined options in each area. The recommended target architecture is as follows:

- For inter-facility networking, the use of the *Fiber Distributed Data Interface* (FDDI) standard and associated technologies over single-mode optical fiber.
- For "backbone" intra-facility networking, the FDDI standard and associated technologies over multi-mode fiber with a 100 million bit per second (Mbps) bandwidth.
- For "access" intra-facility networking, the use of the IEEE 802.3 (Ethernet-based) standard over unshielded twisted-pair (i.e., 10BaseT) with a 10 Mbps bandwidth.
- For workgroup interconnection, the use of intelligent multimedia concentrators, located in building wire centers to interconnect workgroups in conjunction with the access and backbone networks.
- Use of a mixed protocol stack comprised of the TCP/IP suite of protocols and GOSIP-based protocols such as X.400.
- Use of a network operating system (NOS) which provides agency-required functionality as well as "global naming" and low operational costs
- For network management, the use of a "hybrid" approach to include products compliant with the *Simple Network Management Protocol* (SNMP), a NOS-based network management system and special-purpose products that provide supplemental functionality.

Thus, the recommended design or "target architecture" for the agency is a network combining multiple IEEE 802.3 segments over an FDDI backbone.

Multivendor Enterprise Networking (*continued*)

Each segment uses 10BaseT twisted pair cabling to connect individual workstations or servers to intelligent concentrators. Using SNMP, the objective was to manage key components of the network. To address the interoperability issue, multiple LAN architectures were incorporated.

Prototype development

Based on the target architecture, our primary objectives were to engineer, furnish, install, test and evaluate a prototype configuration that demonstrates the functionality needed in the agency's computer and network environment. In addition, the prototype was to demonstrate the performance necessary to meet transmission needs and to provide a testbed to evaluate future networking requirements. To ensure success, we based the prototype on established products, prudent planning and proven vendor experience.

The first tasks in the prototype implementation were the selection, purchase and installation of all necessary components. Due to the aggressive schedule established by the agency, we needed to make these decisions quickly. The criteria used to select the prototype hardware and software included:

- First, due to time constraints, quick delivery was essential. Therefore, off-the-shelf products were necessary.
- Second, with little time allocated for the debugging of hardware and software, we insisted on field-tested products from vendors with good track records on support.
- Third, we required products, as well as vendors to be flexible, since the configuration evolved and products had to accommodate last minute changes.
- Finally, given budget constraints, we selected the best values available in terms of functionality.

We conducted most of our research and analysis at the Interoperability Laboratory headed by Mr. Barry Reinhold at the University of New Hampshire (UNH). Since many of the leading network vendors participate at UNH, critical personnel resources were readily available.

We required two weeks at UNH to build and integrate this system into a working model, and to demonstrate that LANs from different vendors could work together in a multivendor enterprise network. Figure 1 shows the functional prototype as it was constructed.

Tests and results

Once the prototype configuration was certified to be operating according to functional and performance specifications, we deinstalled the system from UNH, transported it to the agency headquarter in Washington, D.C. and reinstalled it for further testing. Although we ran many tests on the network, the following six (6) tests formed the nucleus:

- *Universal access to multiple network operating systems:* This test demonstrated that the target architecture and associated hardware and software allowed users to access LANs using different network operating systems from a single workstation. All accesses occurred across the FDDI-based optical fiber backbone.
- *Moving devices between different locations or buildings:* This test demonstrated the ability to move between different locations while preserving communications with the workstation's original workgroup LAN.

- *SNMP network management:* This test demonstrated the feasibility of using an SNMP-based network management system to view statistics such as the number of packets, collisions, and alarms on a variety of network equipment.
- *Proprietary network management:* We used a proprietary management package for the monitoring of the fiber optic backbone. The test showed the usage of a proprietary system to perform functions necessary to monitor a fiber backbone and the bridges that connect to it. We found that this type of system will be necessary for the short term until standards-based network management systems become more mature and provide a full range of network management features.
- *Use of the TCP/IP protocols across the fiber optic network:* We tested the usage of native-mode and gateway implementations of the TCP/IP suite of protocols across the fiber backbone for file transfer. We were able to move files to and from all network architectures.
- *Electronic mail using the X.400 Message Handling System (MHS):* This test demonstrated the usage of X.400 MHS gateways. In all, we were able to send and receive messages and documents between systems using three different electronic mail packages.

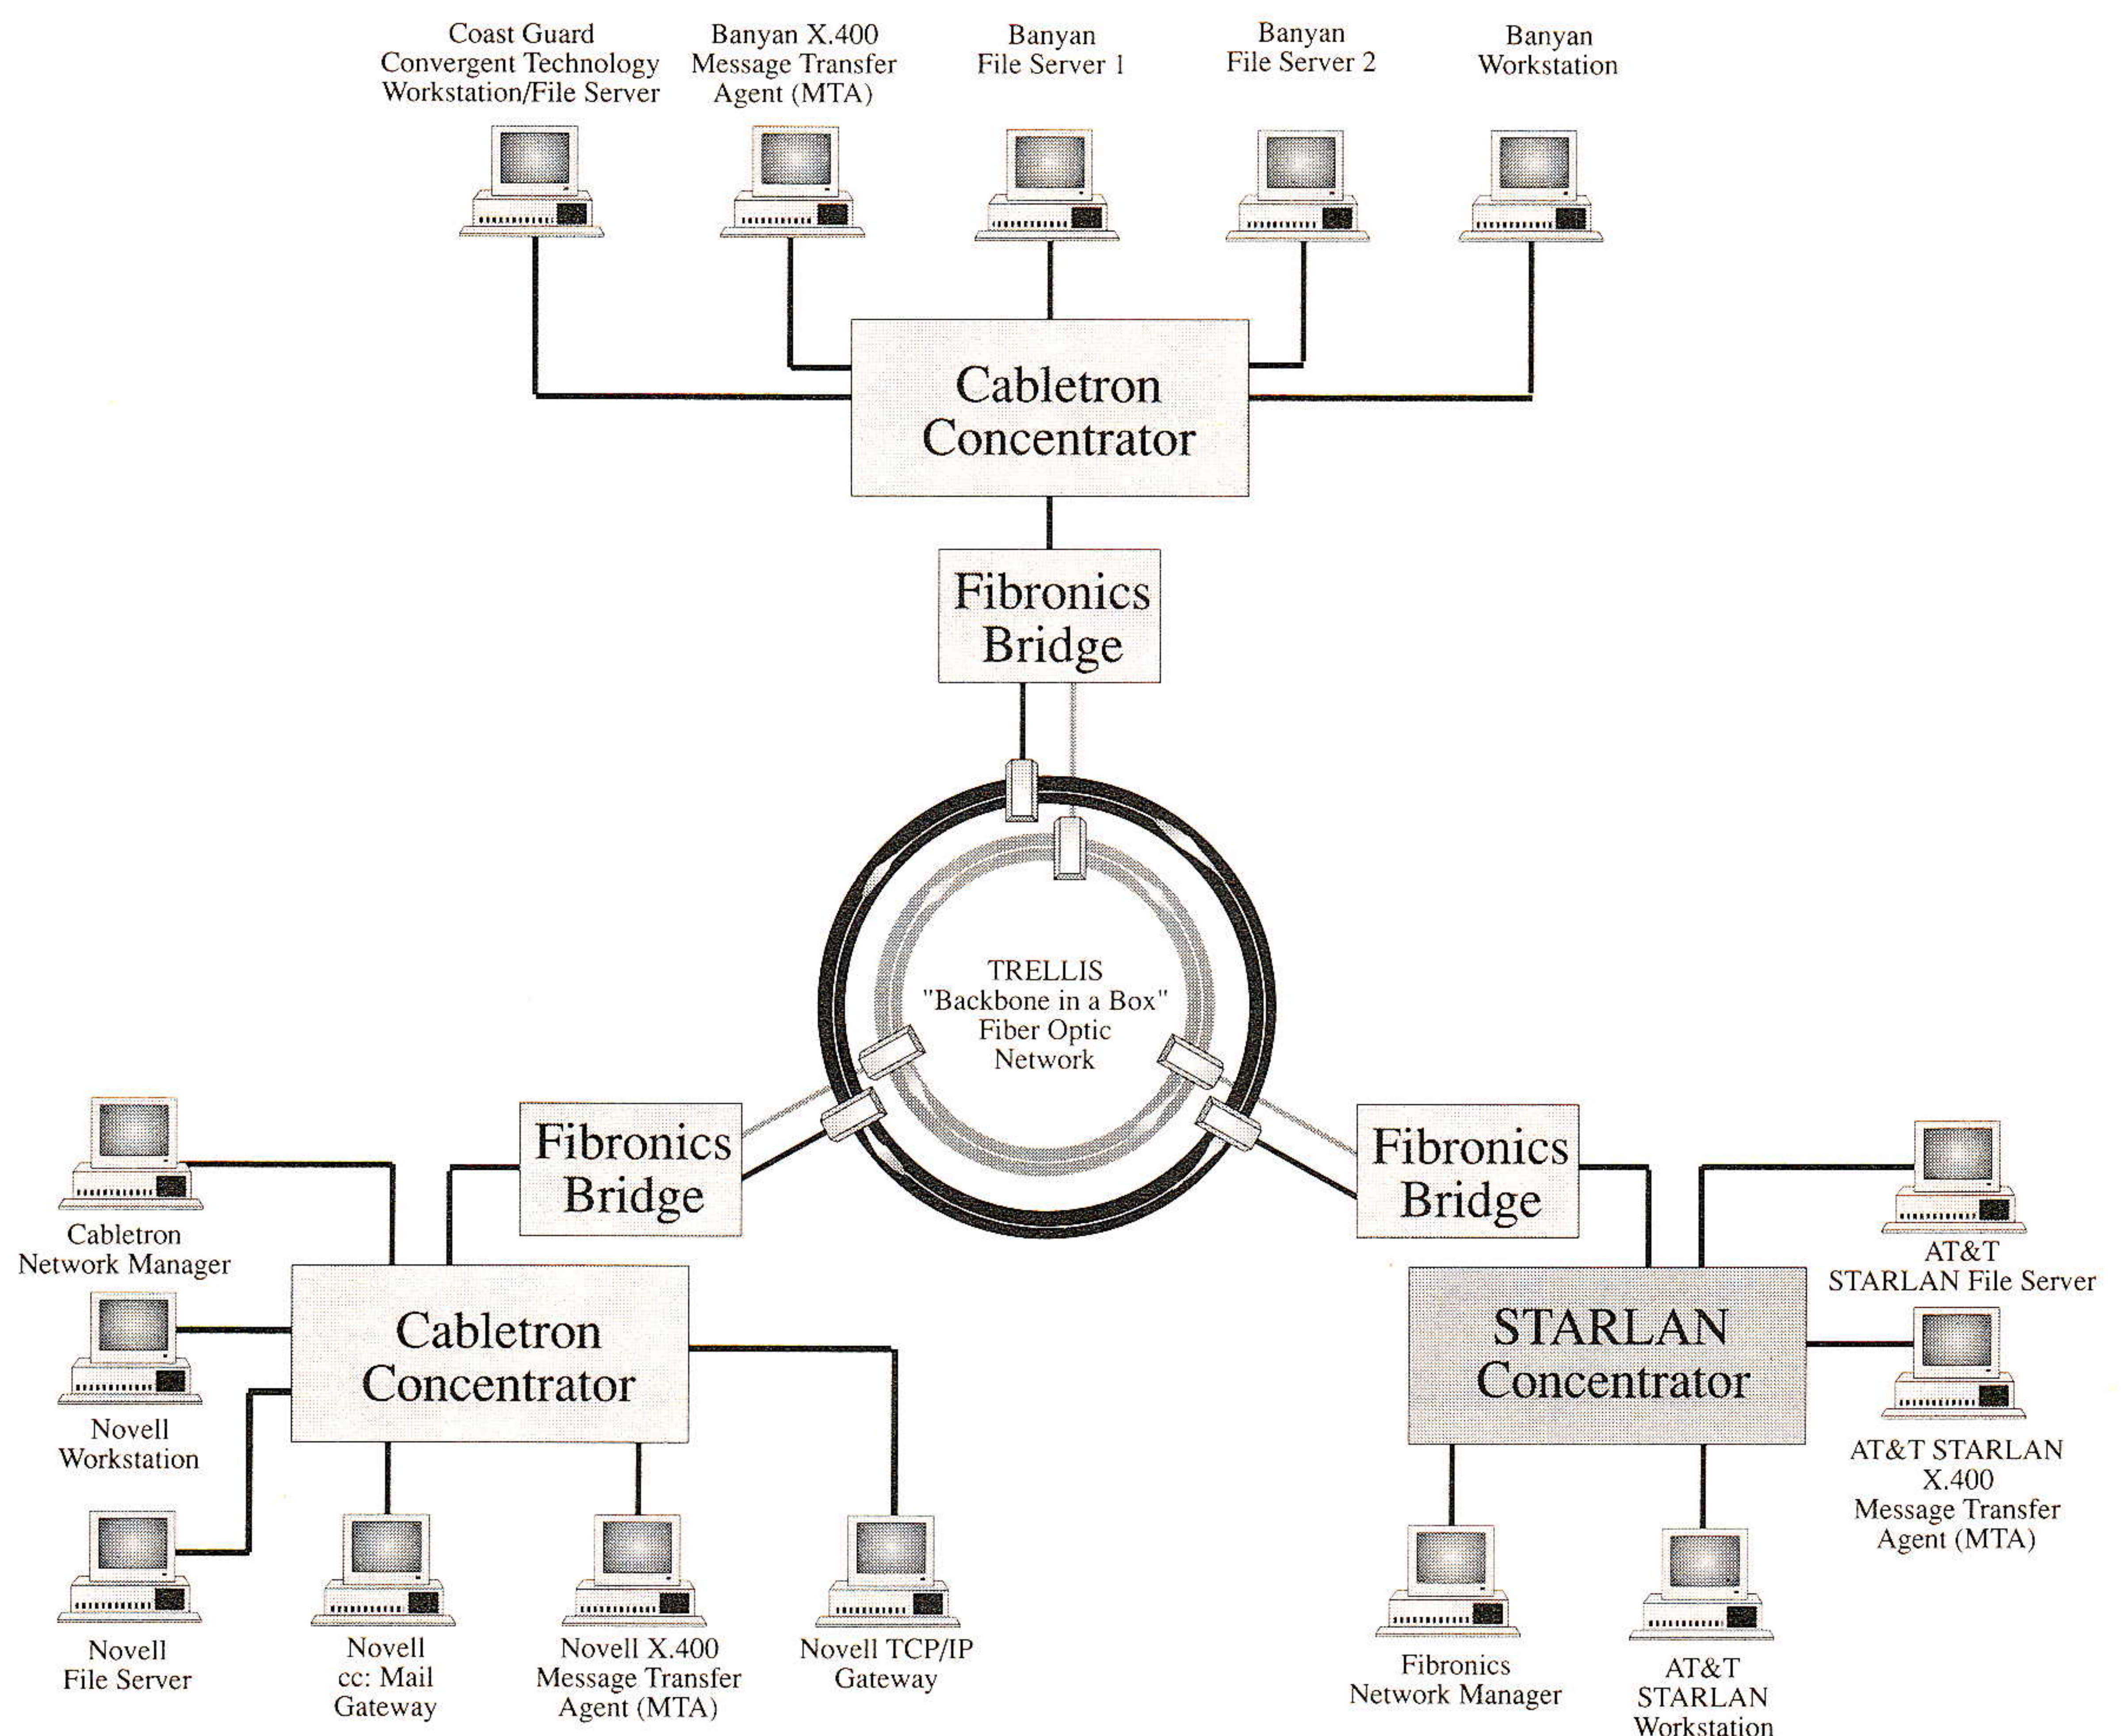


Figure 1: Prototype functional design

continued on next page

Multivendor Enterprise Networking (*continued*)

Lessons learned

The objectives of the prototype were as follows:

- That the target architecture can work in a real world environment such as a government agency;
- That there are hardware and software products available commercially to implement the target architecture; and
- That interoperability at various levels can be achieved among multiple vendors' hardware and software products.

With these objectives in mind, the next paragraphs describe the lessons learned from implementing and operating the prototype. In our prototype demonstration, we found that some products worked as advertised, some did not work and some worked almost as advertised. The following represents the key lessons learned:

In general, the prototype confirmed that the target architecture was implementable. As the first task in implementing the prototype, we conducted a market survey of commercially available products that had the potential to be used to implement the target architecture. We found that there are products available from multiple sources that were field tested. No product selected for the prototype was a test version. However, some vendors do have more mature products, provide better support and assistance and were willing to furnish a corporate commitment to ensure interoperability with other vendors' equipments. However, we found that not all implementation and integration issues could be addressed and resolved during the time-frame when we demonstrated the prototype results.

There are differences among commercially available network operating systems (NOS). We analyzed the major network operating systems three of which (Novell *NetWare*, Banyan *VINES*, and AT&T's *LAN Manager*) were included in the prototype demonstration. We found major differences in performance, functionality, operational costs, and the ability to support global naming.

Standards compliance

There are potential integration pitfalls related to vendor compliance with network standards. We found that there are some shortcomings with vendor compliance to particular network standards. A key finding is that there are currently limitations to being fully GOSIP-compliant. Standards demonstrated in the prototype included the FDDI standard, the X.400 MHS standard and the 10BaseT standards:

- *FDDI*: The FDDI standard is essentially complete and in final form, but vendors have chosen to implement FDDI in slightly different and sometimes incompatible ways. One method is through the use of an "encapsulating" bridge, the mechanism used in the prototype demonstration. With this method, special vendor-specific FDDI-based headers and trailers are appended to Ethernet packets which then traverse the fiber optic backbone network. Using this method precludes the use of other vendors' FDDI products on the backbone. This situation is similar to that which existed when products complying with the Ethernet standard first came into the marketplace in the early 1980s. Then, as now, not all products complied with the standard in exactly the same way. This problem resolved itself in one to two years as the Ethernet standard became more ubiquitous. The FDDI problem should resolve itself in the 1991–1992 timeframe through the use of "translating" bridges.

This type of bridge translates from a given IEEE 802 protocol (e.g., IEEE 802.3 (Ethernet) or IEEE 802.5 (Token Ring)) to the FDDI standard. With this method, vendors products are forced to interoperate with one another in order to compete in today's market.

With the evolution from encapsulation to translation, the government agency will have the ability to attach high-performance CAD workstations with FDDI interface cards directly to the backbone without the use of bridging and at the same time allow other slower speed devices to be bridged onto the backbone.

- *X.400 MHS*: This electronic mail standard is GOSIP-based and has received widespread attention from users who desire to integrate incompatible electronic mail systems within an organization. The X.400 standard is based on the concepts of *Message Transfer Agents* (MTAs) and *User Agents* (UAs). In the prototype, we implemented these concepts through the use of a particular vendor's X.400 MTA product which supports X.400 gateways for proprietary electronic mail systems as well as the vendor's own X.400 "native" mode electronic mail package.

In addition, we attempted to add a UNIX-based and a Macintosh-based mail system to our X.400-based repertoire through the use of a second vendor's MTA product. However, due to slight differences in the implementation of the X.400 standard, we were not able to get the two vendors' MTA products to interoperate. At the time that this article was being prepared, the two vendors, Retix and Touch Communications were working together to determine where the problem areas are.

In addition, one department of the government agency was using Unisys to implement an X.400 gateway for its Convergent Technology Operating System (CTOS)-based electronic mail system. We found that this implementation did, in fact, work with the Retix product. Our prototype showed that three different electronic mail systems (i.e., CTOS-based electronic mail, *cc:Mail* and *RetixMail*) could interoperate using the X.400 MHS standard.

- *10BaseT*: The IEEE 802.3 standard, which implements Ethernet over twisted-pair wire, became a final draft in September 1990. In developing 10BaseT, vendors and standards organizations worked closely to ensure that the specifications were very explicit and that there was little room for customization. As a result, 10BaseT products are on the market and can interoperate. In the prototype, we used both Cabletron and AT&T media concentrators interchangeably.

Problems

Implementing the prototype was relatively easy, but some problems did occur. To implement a system consisting of multiple vendors' hardware and software, it was necessary to have user and systems manuals that provided clear explanations of concepts, procedures and diagnostic capabilities of the products. In addition, qualified support from vendors either through on-site or telephone assistance was a necessity. Most vendors participating in the prototype were responsive and provided excellent support. However, some vendors were not as responsive. The biggest problems were difficulties in (1) finding systems engineers who could give definitive answers on why hardware or software did not work as documented, (2) determining why random failures of system software would occur (some problems still remain unresolved), and (3) finding system engineers with more in-depth knowledge to diagnose interoperability problems with other vendors' products.

continued on next page

Multivendor Enterprise Networking (*continued*)

Performance

From a qualitative perspective, there were differences in performance of the three primary LANs. While the prototype demonstration was essentially functional, there were noticeable differences between the network operating systems in such areas as rebooting and response time when changing applications. Our observations confirmed that the Novell system software is optimized for best overall performance as industry studies indicate. However, we also found that the differences between Banyan and Novell were small when rebooting or changing applications. AT&T's implementation of LAN Manager was slower in these areas. Since all server hardware was equivalent in the prototype, it is likely that the AT&T performance would improve with an upgraded server configuration.

Bridges

"Filtering" and "self learning" Media-Access Control (MAC)-level bridges work well in this government agency's environment comprised of three buildings in a metropolitan environment. All resources that use the IEEE 802 standards have MAC-level addresses (e.g., the 3Com and Cabletron interface cards used in the prototype). With these types of bridges, it is possible to (1) segregate local traffic from department-wide traffic, thereby eliminating potential "broadcast storms" (i.e., the type of problem that occurs when a device mal-functions and sends out a continuous stream of packets onto the network) and (2) move a device from one location to another with the bridge's system software modifying its address tables to "know" where the new location of the device is.

Another benefit of these type of bridges is that they operate in the framework of lower-level IEEE 802 protocols and can transparently support multiple higher-level protocols. This is useful in the government environment where GOSIP-based, TCP/IP-based and other proprietary higher-level protocols are likely to coexist for several years as the transition to a fully GOSIP-compliant network is accomplished. We had an alternative choice of router technology for the prototype. However, routers have particular protocols such as TCP/IP as part of their system software and tend to not offer the flexibility of more intelligent bridge devices.

E-mail

At this time, not all electronic mail systems in the marketplace or within the government agency have gateways to the X.400 MHS. In the prototype, we used three electronic mail systems: (1) cc:Mail, (2) RetixMail and (3) CTOSMail. cc:Mail is one of the most popular and widely used PC-based electronic mail system. As a result, gateways to translate from cc:Mail's proprietary protocol and format to the X.400 MHS standard are commercially available. Retix, which is a supplier of X.400-based products, provides RetixMail as a package that is fully compliant with the standard. CTOSMail, which operates in the closed CTOS environment, has an X.400-based gateway which was developed, tested and shown to interoperate in our prototype environment. During the 1991-1992 timeframe, we expect to see many more X.400 gateways for commercially available electronic mail products.

Network management

The Simple Network Management Protocol (SNMP) is where the industry is moving in network management for a multiple vendor environment. During the timeframe in which we implemented the prototype, the INTEROP 90 conference presented a showcase demonstration comprised of many vendors implementing SNMP products. These products were connected to an FDDI-based fiber optic backbone network and a subset of management functions were performed between the systems in the prototype.

Summary of benefits

In the prototype, we demonstrated SNMP with Cabletron's *LAN-VIEW* product. Future work will be focused on the use of multiple vendors SNMP products as well as the use of the *Common Management Information Protocol* (CMIP) which is expected to be mandatory for the government in the next 2-3 years.

There are a number of benefits associated with the architecture developed for this government agency. First, given the large number of moves and changes of personnel (i.e., approximately a 30% rate per year), there are quantifiable cost benefits to adopting a standard, universal wiring plan based on the FDDI backbone and the 10BaseT access network. Our estimates indicate that the cost savings for an agency of this size would be approximately \$500,000 per year in rewiring charges.

We designed the architecture in several "design areas" for two purposes. First, the total architecture provides a "blueprint" for the agency to aspire to over the next several years. Second, the end user can choose the level of interoperability most applicable to his organization to utilize existing equipment. This was done by developing a "charge back" system and mapping the design areas into various costs to be charged to the end user organization.

For example, if an end user organization has several divisions through a single building and there is already existing Ethernet coaxial cable, the architecture can provide an attachment to the media concentrators located in the nearest wire center which in turn allow the user to bridge to the FDDI backbone. In this scenario, the user only pays for the usage of the fiber optic backbone.

In a second scenario, there are legal and budget administrative counterparts in all departments of the agency. These groups exchange budget information files as well as large amounts of legal documentation. The current method is manual through courier services since electronic mail systems and file transfer protocols are incompatible. With the use of X.400 MHS, these systems can communicate across the fiber backbone. In this case, there is a higher degree of interoperability and thus, the users are charged accordingly for the services used.

A third benefit relates to network management. Through the use of SNMP, multiple types of networks can be managed from a single location. The current situation utilizes much personnel in the network management role with each organization employing different procedures. With this standard, we estimate that the agency will save over \$1,000,000 per year in personnel costs for network management.

There are other benefits, such as the use of easy-to-use, common user interfaces across multiple LAN operating environments and global naming for assistance in configuration management. In summary, the architecture developed for the government agency provides a growth path and allows users to set their own pace in achieving interoperability among their peers and other agencies with whom they communicate.

WAYNE RUSSELL is a telecommunications consultant specializing in large corporate networks for both the government and private sector. In his twenty year career, Mr. Russell has worked in the Naval Research Laboratory community, been a Principal and senior consultant with American Management Systems in Arlington, Virginia and served as an independent consultant with Microage Architects in Fredericksburg, Virginia.

INTEROP 91 Fall will present this case study in the "Emergence of the Common Corporate Internet," session S28, at 3:30 pm on Thursday October 10, and the Corporate Internet BOF the same evening at 9:15 pm.

Networking the Telecom Standards Bodies

by

A. M. Rutkowski, International Telecommunications Union

Introduction

It is a classic example of the cobbler's children who had no shoes. Almost all of the scores of bodies throughout the world engaged in making telecommunication and information standards remain themselves without significant electronic internetworking capability.

Although virtually every standards body has its material in machine-readable form, and many have internal LANs, almost none have external access to that information; and there is no networking among the bodies. All the documents and standards are available only on paper through a few—and often user unfriendly—distribution channels. The only notable exception is the *Internet Engineering Task Force* (IETF) where everything is coordinated and openly accessible through the global Internet.

The standards distribution problem partially arises from attempts to maintain artificially high paper-copy prices by creating a monopoly through a copyright assertion—a practice effectively impossible with open electronic internetworks. However, almost anyone with one the new scanners available can easily provide an “optical gateway” into the network domain. If a sizable black market in electronic copies of standards does emerge, the standards bodies will not be able to put the electronic genie back in the bottle. It is important for the bodies themselves to provide what users are demanding.

Changes

Some major changes are in the wind, encouraged by several factors. These include: rapidly changing technology and markets; participatory costs and lost expertise; new global open market norms; and increasing competition among the standards bodies.

Perhaps most significant is the very recent availability of the needed internetworking capabilities on a scale that makes it feasible to apply to the standards bodies worldwide.

This article was prepared specially to stimulate wide robust discussion on this important subject among all the affected professional communities and standards bodies. It reviews the pertinent factors and recent developments, suggests how telecom standards inter-networking could now be achieved almost immediately, and why it is in everyone's common interest to do so.

Networking tools

During the past 2 years, the information networking world has witnessed a revolution that is profoundly changing how organizations, professions, and individuals share information and collaborate in their work. The revolution is centered around the interconnection of thousands of information networks around the globe to form the open, multiprotocol, cooperative metanetwork called the *Internet*.

The Internet began growing rapidly in the U.S. in the late 80s—more than doubling in size every year. Over the past two years, that exponential growth pattern was replicated around the globe. Most of the telecommunications world is now connected; and shortly most of the remaining geographical world will be. Current users exceed 3 million and are expected to reach 300 million by the mid-90s.

Connectivity and interworking are very simple and inexpensive through modems, local area networks, multiprotocol routers, and registration. Anyone can connect. The architect of the pan-Australian Internet backbone network AARNet recently described it as “like going to K-Mart.”

There are now scores of public initiatives and commercial ventures around the world to build and operate national and regional networks that are part of a global Internet. With access comes several simple, basic tools that include the ability:

- To exchange messages with millions of users,
- To search through and transfer files from thousands of open information hosts,
- To access supercomputer resources,
- To automatically propagate and receive news on hundreds of specialized subjects.

Non-commercial and research users have free use of these tools since the technologies used are extraordinarily efficient, and because so many national and regional research and development initiatives worldwide all share the costs of the backbone transportation networks. New commercial ventures are now providing the same capacity at minimal cost to everyone. All of this is not the future—it is today.

Current users

It is ironic that although the telecom standards bodies are not using internetworking tools, dozens of other communities ranging from high-energy physicists to primary school children are using these tools as a natural, important part of their daily environment. Indeed, I find that when almost anyone in college today visits the International Telecommunications Union (ITU) in Geneva, they inevitably find their way to my office to Telnet (remotely log in) to their home host computer.

Large active communities of diverse professionals now share information and collaborate around the world. It has even spawned entirely new disciplines like collaboration theory and resource discovery.

For example, those in high energy nuclear physics distribute bursts of experimental data from collision experiments for near real-time analysis and share supercomputers. This allows a physicist in Australia or China access to most of the same data and resources as if she were at the accelerator site.

Molecular biologists share information on complex protein molecules and contribute to the massive on-line data bases that automatically swap information between Europe and the US for mapping human genes. Most of the world's major library catalogues are accessible, and librarians are collaborating internationally to establish standard automatic subject search capabilities. Automobile designers are sharing design ideas with their brethren in other countries. Physicians are forming speciality groups and using on-line medical references at the best disease centers. The list of collaborative communities today is pretty big. But it is young people who seem to really love these tools.

Educators are establishing global specialty groups—even projects to create “The Global Classroom.” For example, a poor inner-city Hispanic primary school in Boston is internetted with a rural school in Costa Rica, allowing the children to share messages and files of video snapshots and drawings of their environment. Young students in Prague are forming discussion groups with counterparts in Australia and the US. Psychologists in the Soviet Union are participating in collaboration theory discussions with counterparts in San Diego, California.

Networking the Standards Bodies (*continued*)

And to make it even easier, the same MIT folks that brought us The X Window System are putting the finishing touches on *LogoExpress* to allow the five-year old crowd to internetwork. Waiting for them is *KidsNet*—formed out of Norway to emulate a café scene for children around the world to meet.

Applying the tools for standards making

Those engaged in telecommunication-information standards making today have needs that are almost identical to most other professional communities. In fact, the highly distributed and autonomous architecture of the standards world today (shown in the attached chart on page 35) is the very image of a distributed network. Every one of the bodies have three basic needs:

- Fostering collaboration (presently largely in the form of meetings and via messages).
- Redistributing information (material provided to or solicited by the standards body that is subsequently redistributed, e.g., meeting documents, questionnaire answers, chairpersons and rapporteurs, group lists, etc).
- Distributing information (internally generated notices, news releases, or standards).

In performing these tasks, no standards body today stands alone, but is already part of a complex, increasingly non-hierarchical matrix of bodies where the information is constantly being transferred, compiled, and adapted among hundreds of different organizations.

Already many of the companies and individuals that participate in standards making activities are part of the Internet, and increasingly in CCITT and IFIP groups, Internet SMTP mail address are regularly found next to people's names on the documents. If a poll were taken today, it would likely show that most companies participating in telecom standards bodies have either direct access or gateway e-mail access to the Internet.

Every telecommunication-information standards body in the world is near enough to an existing Internet node, that with the simple addition of a short local leased line and a multiprotocol router, connectivity among and with every one of those standards bodies could be attained.

With just the basic tools of mail, file transfer, and remote log-in, the benefits to everyone associated with the global standards making, manufacturing, and service provisioning communities would be enormous—and it could be done within a month!

Why telecom standards bodies should be internetworked

The fact that the tools exist for internally and externally internetworking the standards bodies will not by itself compel them to use the tools. However, there are many other developments that should provide strong motivation.

- *Rapidly changing technology and markets:* At a recent meeting of a major standards organization, a manufacturer delivered an eloquent message on the microphone. He simply said that the information-telecommunication technologies and markets were changing so fast today that his company had only about 18 months from the initial feasibility of a product offering to its release, and that if a standard could not be developed within that timeframe, it couldn't be used. And even then, he noted, each additional month represented major lost opportunity costs.

With relatively few exceptions, the 18 month rule is the norm in today's information systems world. Even then, it is often necessary to adjust specifications to comport with constantly advancing capabilities in basic technology implementations in processors, memory, and transmission speeds.

- *Participatory costs and lost experience:* The costs of participating in standards making activities have risen dramatically. These costs are not only actual cash layouts in terms of salary and travel to attend the ever growing numbers of meetings, and reviewing and writing documents. Also significant are the costs of losing expert individuals for weeks at a time around the year to meetings which often use their skills inefficiently. Most companies are finding it increasingly difficult to support such costs; and the results are reflected in the current attendance lists of many standards meetings—where participation is too often skewed in the direction of large players and particular industry sectors.

Smaller companies and academic institutions—where many of the most creative and “hands-on” users of the technology abound—are effectively shut out of most of today's traditional larger standards making activities. Also effectively excluded are individuals from resource-limited developing countries. It is very difficult to get the documents which are almost exclusively distributed at the meetings. It is very difficult and exceptionally costly to get current standards in draft or even final form.

Copyright

Distribution is further impeded by copyright assertions of many standards bodies. Again the only exceptions are the IETF and new “startup” industry-user standards forums which are developing many of the industry's most successful standards by becoming meccas for innovative and enterprising individuals and companies.

- *Widespread pirating:* The cost and copyright concerns have already led to widespread photocopying of standards. It has become almost equally easy with good, inexpensive, optical scanners and software to convert paper copies back into electronic images and place them on a server. Doing this is getting cheaper and easier by the month. There are a few locations already making available some unauthorized electronic versions of standards.

If enough counter-culture people, research bodies, or even some “dare to sue me” commercial ventures scale up these activities, the standards bodies will lose control over the electronic standards distribution process.

It is eminently more sensible for the standards bodies themselves to recognize the need for robust, widespread distribution of good electronic copies of standards—and the resulting benefits to their organizations, the engineering profession and the industry.

- *Competition among the Standards Bodies:* One of the most significant developments in the telecom-information standards making world is the tremendous growth in the number of bodies engaged in this activity. It is, as CCITT Director Theo Irmer often repeats, “a competitive business.”

No one would argue that more standards are not needed in today's rapid paced, interoperating digital world. On the other hand, many of the new bodies have arisen for reasons other than just making more standards.

Networking the Standards Bodies (*continued*)

Layering effect

One major factor is that no one body can do all the work in the required timeframes with the required specificity with the necessary service to local constituents. As a result, there is a layering effect where global bodies like CCITT and ISO make general—often abstract—standards or models with many options that are not capable of singular implementation. Because this lack of specificity often arises from disagreements among participants in the work, it is possible that earlier, faster, less formal electronic collaboration among the participants might bring about more complete standards at the global level.

Often, these standards have never even been tested to see if they actually work—even as they are adopted. (The IETF is perhaps unique in explicitly requiring extensive testing of a draft standard prior to adoption). It is usually left to regional and/or national bodies and/or individual providers to then develop, flesh out, and test more detailed, implementable standards. Even separate conformance testing standards bodies have sprung into existence over the past few years to focus on the testing requirements alone.

The quandary is that traditional standards making has tended to become a big, complex, and often inefficient business in an era where the market demands efficiency and speed. Recent new fast-track approval procedures pursued by some bodies like the CCITT and CCIR have significantly accelerated adoption timeframes.

Still, significant liabilities remain because many of the big standards bodies fall prey to the tendency of all big institutions to devote large amounts of resources and energies to overhead. Such overhead includes costly activities devoted only to institutional needs—especially needless rote translations and inter/intra-body liaison paper—as well as the pursuit of standards that the ultimate potential consumers of standards don't need, perhaps don't even want, but serve the interests of a particular participant. It is the Sorcerer's Apprentice syndrome at work. Good electronic internetworking tends to promote better management, maximize horizontal communication and minimize needless ritual transiting through hierarchies.

Duplication of effort

Another aspect of the same problem is the duplication of effort that takes place among all the standards bodies simply because of a lack of knowledge that someone else is working on or even has completed similar standards. Sharing information resources, including project management data, through an internet—perhaps even combined with automatic search capabilities—could save enormous monies and time, and also result in more globally uniform implementations for the same system or feature.

Another major often-ignored factor is simply that the information-telecommunication industry has become very much more heterogeneous. New generations of entrepreneurs from Silicon Valley, Boston's Beltway, or France's INRIA who are interested in implementing virtual reality are not going to fit into the same institution with those who have engineered monopoly public carrier systems for fifty years. There is probably not a basis to communicate, much less work together. So there are inevitably going to be multiple standards making institutions, because institutions are as much a home for cultures as they are for subject matter.

However, in today's digital world, the subject matter significantly overlaps, and it is going to be increasingly critical to bridge these "cultural" barriers by allowing those in different institutional homes to collaborate.

For all these reasons, there are today many standards making forums effectively in competition with each other, and it is the marketplace, not the status of the institution which will largely decide which standards products are used and which are not. It is the users and the industry itself that are the benefactors because that is what the competitive marketplace is all about—maximising the benefits to ultimate consumers of limited resources in a real world.

- *New Global Open Market norms and initiatives:* Layered on top of all other considerations are emerging new regional and global policy requirements based on antitrust and trade principles and reflected in legally binding agreements and law. The most notable are the draft *General Agreement on Tariffs and Trade* (GATT) framework for a *General Agreement on Services* (GNS) currently being negotiated, and the various "open network" regimes promulgated in the European Community, the United States, Japan, and Korea.

All of these developments require standards making processes be transparent and open, provide prior notice, and easy access to the resulting standards and drafts.

In the GATT, a widespread consensus is emerging in the direction of fair open global markets in equipment and services among the participating governments; and standards making has received close scrutiny because standards and testing requirements although generally beneficial, can also be used as de facto trade barriers and market impediments.

In addition, the so-called *Standards Summit* process initiated at Fredericksburg, Virginia, in 1990 and subsequently continued as an inter-regional conference, is potentially leading toward an increasing cooperation among all standards making bodies in their own mutual interests in achieving much greater collective efficiencies. If this process is really opened up to all organizations, it could provide a common institutional platform for global networking of telecom standards making bodies.

Overcoming the obstacles

Arrayed against all of these motivations toward networking are several serious obstacles—many of which are not very candidly discussed. These include revenue and product control, protocol wars, security fears, put it off until there is something more elegant, and local cultures.

Revenue and product control

Many—but by no means all—standards bodies or their agents assert a copyright for their standards. The practice is usually justified purely as a pecuniary measure, but some arguments for "control" over the distributed copies are also expressed. Other standards bodies—and virtually every standards participant and user—note instead that the test of success of a standards body is the extent of *implementation* of its standard, and not the income derived from the sale of documents. And although the counter argument raises the issue of income to these bodies, most of them exist for and consume enormous participant monies—on the order of several hundred million dollars a year—for the single purpose of producing standards that are actually needed and used. These participant expenditures far exceed revenue derived from sales of standards.

Networking the Standards Bodies (*continued*)

Ironically, many organizations have found that the markets for paper-based and for electronic copies of standards are very distinct. In other words, the market for paper-based versions does not significantly diminish when electronic copies are made available.

Despite the major broader interests at stake in standards making, it is the pursuit of revenue obtained from the sales of paper copies of standards alone that impedes most standards bodies from electronically disseminating their standards—a position that seems badly balanced at best. Certainly the unquestionable great success of the IETF in getting its standards implemented around the world is in part due to the ease with which anyone can retrieve the standards via anonymous FTP out of network servers.

The concern over control of the accuracy of standards can be handled through a combination of easy availability from authoritative servers and simple optional licensing or registration schemes. Internet communities already use these techniques for important reference material. For example, those who obtain a copy of a standard can be explicitly licensed to use the copy for their own purposes only; and users can optionally register to obtain subsequent versions.

Protocol wars

One of the more contentious subjects already resolved by the inter-networking community is the issue of competing protocols—often referred to as “religious wars.” Different factions favour their own pet protocols—TCP/IP, DECnet, SNA, AppleTalk, OSI, etc—particularly factions which develop them for their own competitive advantage or someone else’s disadvantage.

The problems arise when the factions argue (always for the most noble reasons) for exclusive use of one protocol, or for gateways or intermediate agents that favour one protocol or transit path over another. After dealing with these battles for several years, the Internet community established the maxim of “give the user an equal choice.” Thus multiprotocol routers and parallel applications are now the norm. Everyone has pretty much agreed that it is connectivity and facilitating use that is important, not protocol wars.

This subject may be harder for some of the standards making community to deal with, because it is their own standards or constituents that are involved. For example, it is tempting for standards bodies that cater to public carriers to require all communications be routed through their own (generally high-priced) facilities and to promote their own standards. Thus X.400 traffic garners a lot of revenue, while SMTP mail does not. There is no reason why both should not be equally supported and allow users their choice based on convenience and cost.

Most standards bodies are probably mature enough to avoid this kind of favoritism in order to promote the common good of the standards body itself, if not the engineering profession and public at large. No public carrier is likely to actually need the extra traffic to survive. The standards making community will be the ultimate benefactor of communications that provide inexpensive collaboration, and dissemination of news and standards.

Security fears

Sometimes organizations are reluctant to connect to an open internet because of the fear of security breaches into privileged files or harm to their local networks. Although these are possible problems, there are simple steps that can be taken that provide effective remedies. An isolated public (anonymously accessible) server is but one example.

**Waiting for something
more elegant**

Many of the largest companies and research establishments—even military facilities—are connected to the Internet. There is no standards body that deals with such sensitive material that it would prevent interconnection.

Another common argument against interconnection is that there are more elegant document standards and approaches on the horizon, and that the institution should wait until those approaches are widely implemented. Of course, there is always going to be some more elegant solution on the horizon. Users, however, generally do not want elegant solutions, they want minimal tools to get the job done. In this case, they simply want access to the electronic file versions which generated the paper usually sent through the post or distributed at meetings.

It is common on many servers on the Internet today for the same document to exist in several different versions: plain ASCII, native format, *PostScript*, and compressed versions of all three. The user simply accesses the directory and transfers the version that suits her or him. Again, it comports with the maxim “let the user decide.”

Local cultures

Lastly, there is the issue of local computer environments that have their own favorite approach—whether it be applications, operating systems, or information agents—to fulfilling the needs of the organization. To a greater or lesser degree, this problem exists everywhere, because people and organizations tend to become familiar with their own self-learned solutions to local needs.

Local cultures need not be a barrier to internetworking. Indeed, the very concept of internetworking was fashioned to accommodate the great diversity of local cultures, machines, and systems that exist. The only common element is at the point of interconnection where common protocols are supported by everyone in the common interest of achieving a metanetwork.

**A mutual interest
to act now**

The global telecom standards community stands at a fairly unique confluence of events where internetworking could be greatly facilitated. The time to act is now.

Some standards bodies are currently experimenting with PSTN and PSPDN dial-in videotext bulletin boards, and messaging using X.400 and private e-mail services and gateways. On an experimental basis, ITU Secretary-General Tarjanne has recently taken an important innovative step by allowing the *Digital Resource Institute* project at the University of Colorado to place standards of the ITU's International Consultative Committees (CCITT and CCIR) on servers connected to the Internet. Usage patterns will be monitored and intelligent directory programs will be tested. The latter would allow a user to request, for example, all standards that deal with <subject A> and <subject B>. (Project descriptions and papers of the Institute are available by anonymous FTP from latour.boulder.edu or by e-mail to schwartz@latour.colorado.edu or carl@malamud.com.) With internetworking, the same kind of automated searching techniques could be easily applied across all the standards organizations—a capability of enormous potential worldwide benefit and cost savings—as McGill University's *Archie* system has already demonstrated together with linked servers in Finland and New Zealand. *Archie* contains a central database for information about archive sites on the Internet and can be automatically searched. (Additional information is available by e-mail to archie@cs.mcgill.ca).

Networking the Standards Bodies *(continued)*

The second session of the Standards Summit of regional and international standards bodies will meet in September at Nice. This is the perfect opportunity for an initiative where all of the many telecom and information standards bodies, conformance testing bodies, information object registration authorities, and industry-user standards forums agree to cooperate in connecting to the Internet and in sharing all basic management information and standards documents.

CCITT's Resolution 18 Group on working methods is meeting in late October in Geneva to lay out user information system support needs. Already there has been a significant focus on providing the CCITT standards making community with significantly greater network tools; and if these tools can be made available now, the opportunity exists to integrate them into the future working methods of the body.

Even as Internet connectivity goes forward, there will be new challenges. It will require, for example, more coordinated management among all the standards bodies to learn how to best horizontally collaborate. Experimentation will be necessary in how to structure and manage complex standard development projects; to encourage shared, timely goals among strategically competitive participants; to examine if, when, and how to standardize; and to optimize information flows through the internet architecture.

At present, the global standards making architecture is fragmented into many isolated camps based on history, membership, and cultures. The barriers to cooperation must be bridged, and the cherished views of institutional superiorities must be diminished. A new kind of "standards democracy" must emerge which compares and supports standards on their merits, and doesn't automatically regard one standards body as intrinsically better than another.

In many cases, more effective administration and project management capabilities need to be developed within many standards bodies. And along with culture goes the necessary workarounds to deal with the many individuals who are reluctant to use electronic information and internetworking tools.

But these are welcome challenges in the face of the enormous benefits to internetworking the standards bodies. With several hundred million dollars a year being collectively invested in information/telecommunication standards activities, the potential monetary savings alone are enormous—not to mention the value of developing standards that are more appropriate, better, and more used.

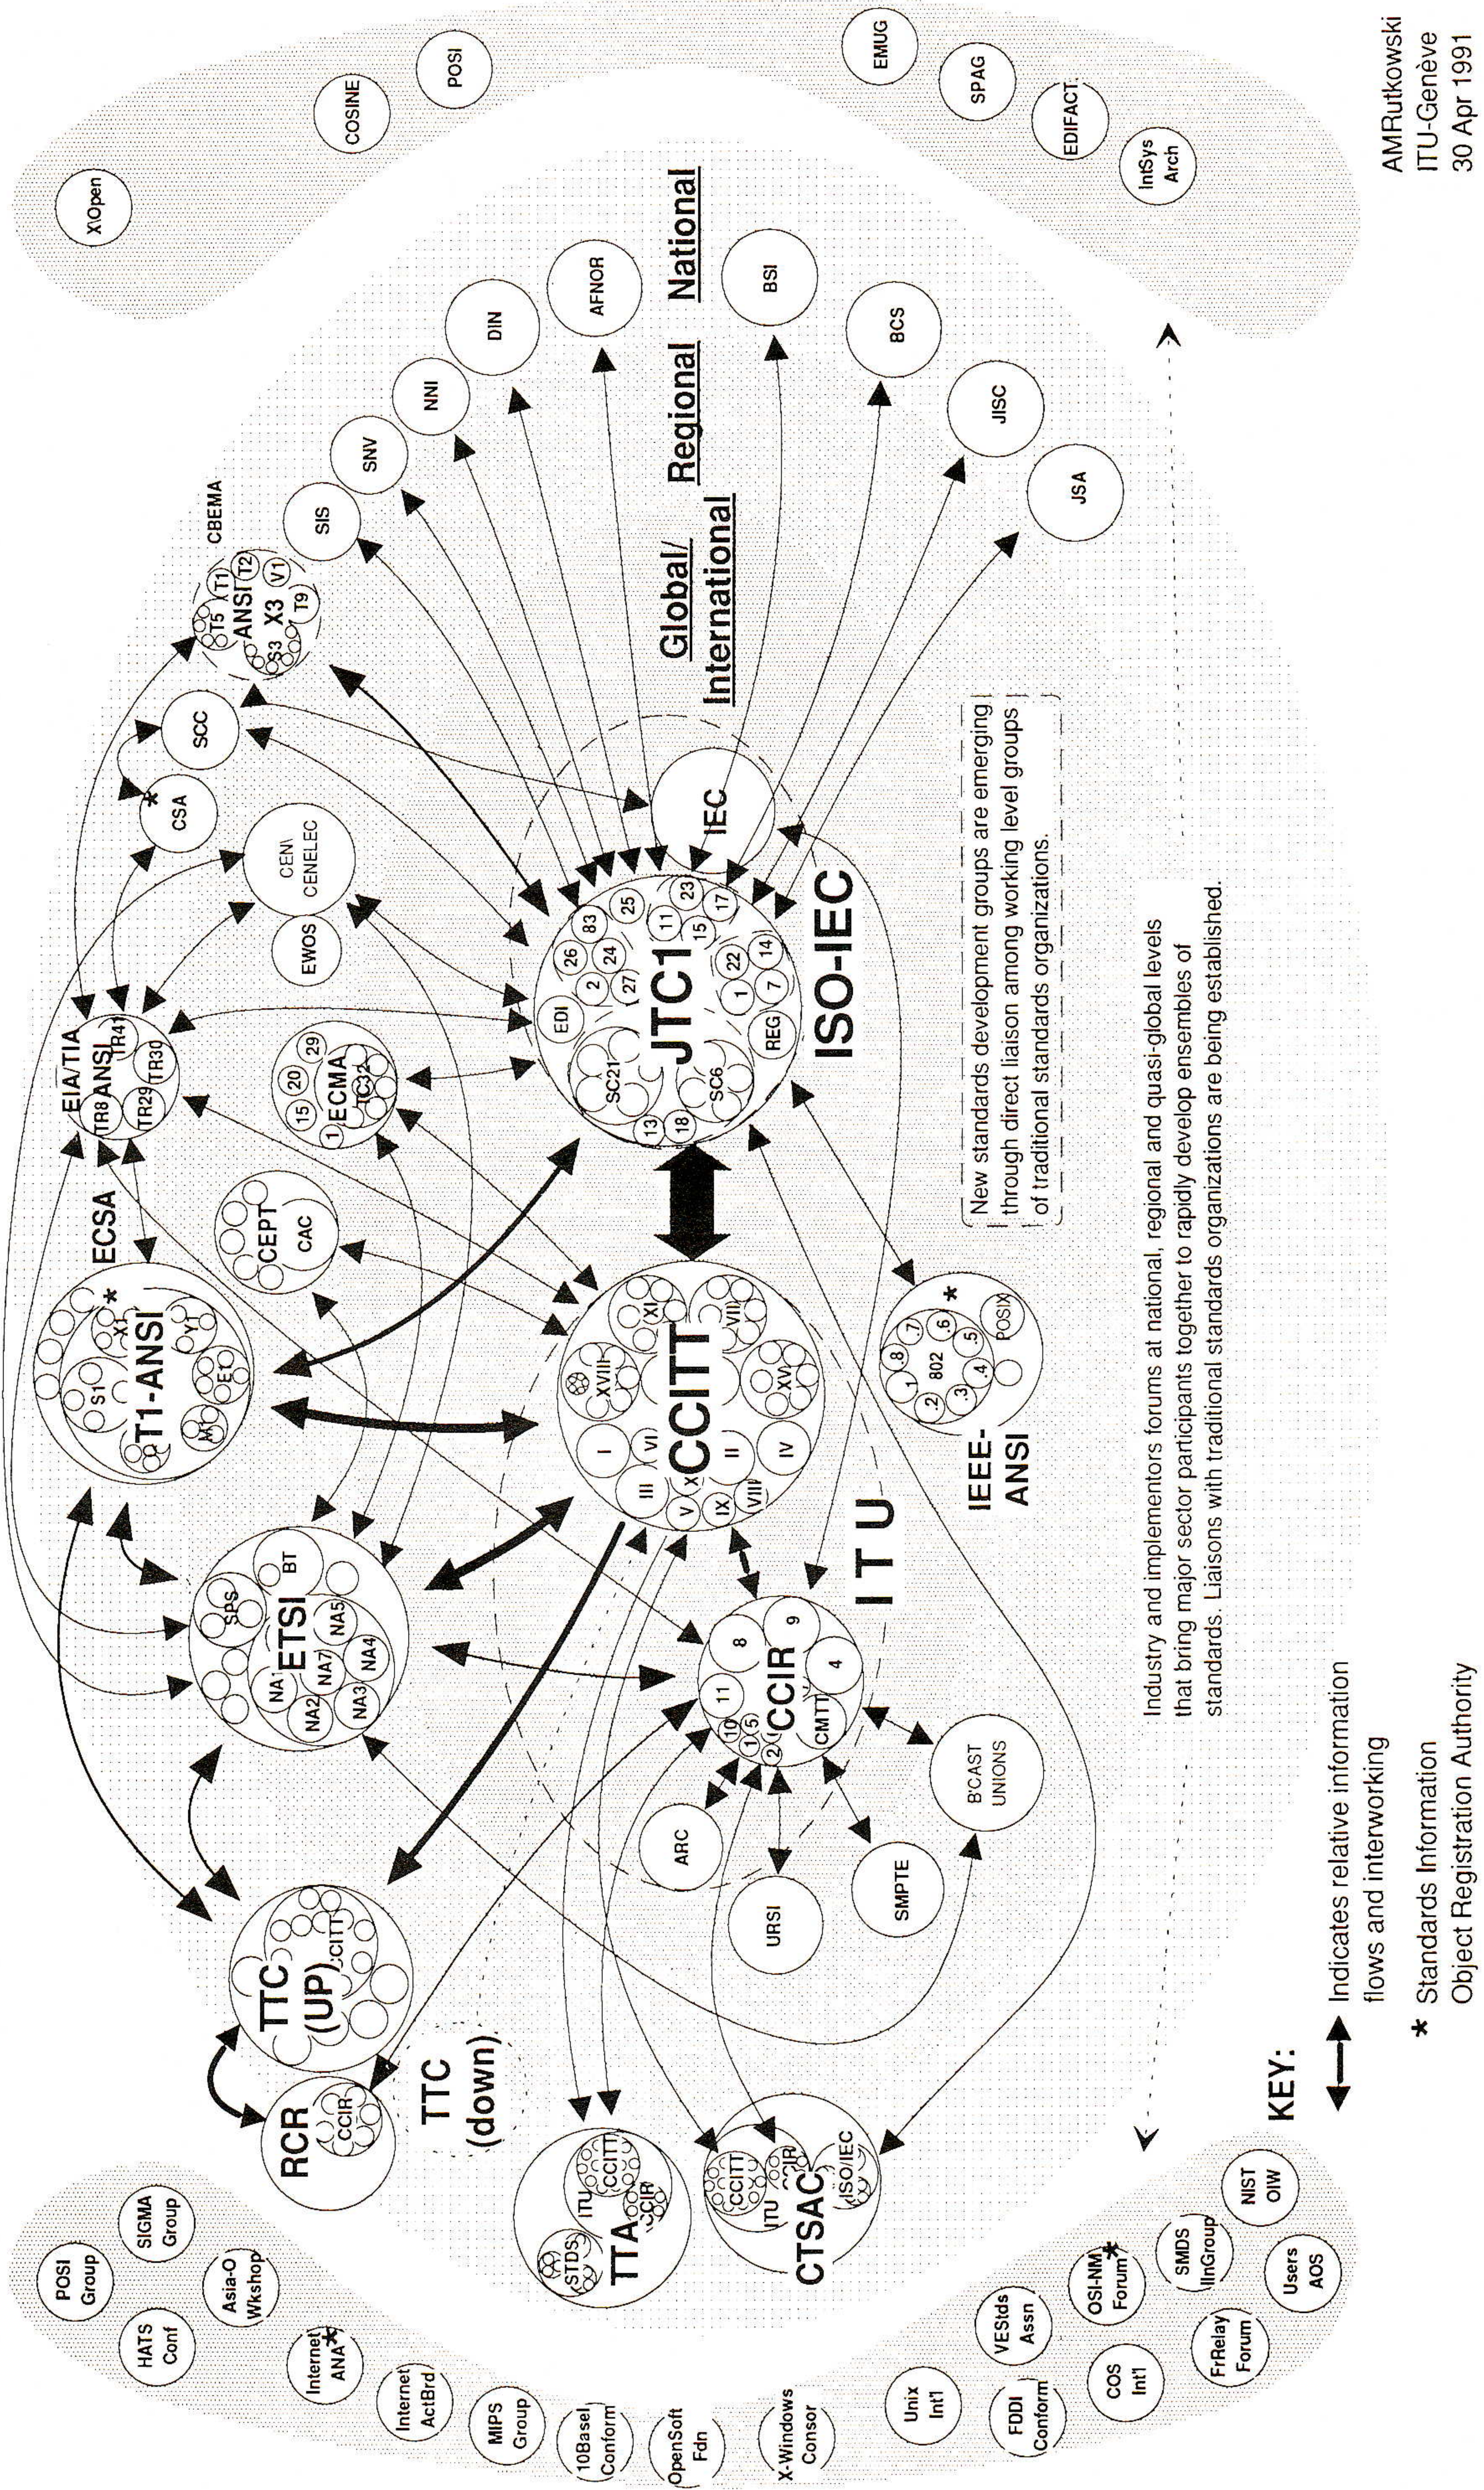
Conclusion

It is hard to imagine today a global community more appropriate for internetwork resource sharing and collaboration than the many telecom-information standards bodies, participants, and users at national regional and international levels. Everyone associated with this community would reap significant benefits. Indeed, such a result will become almost imperative if high level policies envisioning open markets are to be implemented in this economic sector. The internetworking tools to achieve this result are now readily available and easily implemented at negligible cost. It really is time to act now.

ANTHONY M. RUTKOWSKI is counselor to the Secretary-General, International Telecommunications Union, Geneva, Switzerland, and a research associate at the Massachusetts Institute of Technology. The opinions expressed are solely his personal views and are not official positions of the ITU or MIT. Thanks is given to colleagues in several different standards bodies who reviewed and provided additional ideas for this article. He may be reached as amr@cernvax.cern.ch.

Hear more about these important issues at INTEROP 91 Fall. Attend S35 "Who Owns standards?" at 8:30am on Friday, October 11.

Today's Information-Telecommunications Standards Making Architecture



AMRutkowski
ITU-Genève
30 Apr 1991

Building the INTEROP 90 Show Network

by Stev Knowles, FTP Software, Inc.

Introduction

A lot of people had ideas about what they wanted from INTEROP 90. Some people were content to break attendance records. Some folks wanted to make INTEROP 90 unique, offering the attendees box lunches, and the opportunity to go outside for a relaxed lunch to enjoy the beautiful Northern California weather. Some of us were just aiming at making our lives easier. Several groups of people seemed to be conspiring to make the INTEROP 90 shownet something that would be impossible to build, difficult to maintain, and painful to remove. But I am sure you have probably found that these same features seem to encroach on every complex endeavor. As with previous years, INTEROP 90 brought out the best in people, and the worst in people. I won't generalize and say that INTEROP 90 brought out the worst in only specific groups of people, either by job classification, or by corporate alliance. Fortunately, as in previous years, the people who rose to the challenge in a positive fashion far outnumbered the folks who didn't.

Checking the diagrams

INTEROP 90 began for the network folks on January 3rd, 1990. This was the day that Peter DeVries, Dave Bridgham, and myself journeyed to the San Jose Convention Center to do our initial walk-through to get a feel for the magnitude of our job, and to check the accuracy of some building drawings we had been given. This was a useful thing, since the drawings one gets for a building are always either incomplete or incorrect in significant ways. The diagram we had was indeed incorrect in several specific ways, but since we were able to identify these problems early, we were able to work with the convention center in getting the problems addressed.

After our tour of the convention center, which included sticking our noses in all sorts of small, inaccessible areas, looking for places to run cables, we headed back to the Interop Inc. offices. Interop is a great place to hold meetings, they have all sorts of soda and junk food to ingest while pondering how badly you are getting yourself into trouble. At this point, we had 210 booths, 12 hours to hang the wires, and some rough ideas on what to do. I suppose it is fair to let you in on something you will figure out for your self later on anyway: Things were going to get much worse before they got appreciably better.

Cable layout

After several hours of discussions, which mostly involved selecting and refining some of the rough ideas we had, we decided on a modified "rib-cage" design. The convention floor is laid out with the booths in rows going north to south, with one large aisle cutting east to west, just behind the Interop, Inc. booth. Approximately above the east-west aisle is a row of *eye-bolts* hanging from the building steel, sticking down through the drop ceiling. The eye-bolts are mounted in the ceiling in a grid pattern, with the rows doing a reasonable job of following right under the north-south aisles. We decided to run "ribs" in a north-south direction providing network services to the booths. We used 10BaseT wiring, in essence providing the entire row with one small LAN or subnet for their own use. We also ran a bundle down the east-west aisle, with wires originating at each row, and terminating at the middle of the convention floor, where the Interop, Inc. booth happened to be located. This provided us with interconnection points where the north-south rows crossed the east-west spine. At these locations we had small shelves, or "pedestals," as Peter called them, built to hold all the interconnection equipment.

Topology

In the average pedestal was a 10BaseT wire center, a router, bridges, 10BaseT transceivers, and a collection of other random Ethernet equipment, some of which was intended to “spy” on the ribs, to allow us to trouble-shoot problems on the ribs from the show NOC (Network Operations Center). While the ribs had one, or sometimes two, LANs on each segment, the spine carried a large collection of wires into every pedestal. Included were 3 Ethernet backbones, one FDDI backbone, dedicated wires to plug into each LAN to provide people in the NOC with the ability to access any network directly, and wires to connect the console ports for the routers. While this seems like a lot of work to go through, we were very concerned about the manageability of our network.

We were also concerned about the ability for us to access all of our infrastructure equipment in times of total network failure. Each router had a line going from its console port to a terminal concentrator on the NOC Ethernet. This allowed us to be sure we could access any router, no matter what the state of the spine was, as long as our small, internal Ethernet worked correctly. One of the reasons we were so concerned was our desire to show interoperability between different vendors equipment. This could introduce some problems into the shownet that someone with a network of all one vendors products may not have. On the other hand, we expect that more and more customers will have networks like this. If we, “the professionals,” were concerned about this so much that we wouldn’t mix vendors, we really couldn’t ask our customers to do so now, could we? As it turns out, we were fortunate to have built so much control into the network, since we would exercise this control over and over as the tradeshow part of INTEROP 90 approached.

FDDI problems

As I am sure you will know if you have read the trade rags, the FDDI backbone did not function reliably enough for us to feel comfortable using it as our primary backbone. We had 24 routers on the shownet backbone, but found that we could not get more than 18 onto the ring at one time before the ring became unstable. While this could have been a disaster, we had prepared before hand, and had 2 backup Ethernet backbones ready to use in case we had problems. When the FDDI ring was running with 18 routers, we had routers from 4 vendors running at the same time, without problems.

There are certainly people who feel that we pushed too far in trying to get new technology into the forefront in INTEROP 90. If I had to do it again, I would certainly try to include new technology in the net, with the older technology there as a backup. I have also heard rumors that it was “one vendor’s fault” that the shownet backbone did not work. I can state with the authority of the person who was in charge of the repeated attempts to bring the FDDI spine up that *no single vendor* was any more guilty than any other. This was a newly introduced technology, with no clear “lead” vendor to provide guidance. As a person who lived through the trials of 802.5/IBM Token Ring, I can tell you that specs can be read by two people, who can correctly interpret the documents, yet produce two incompatible implementations. The ability to point to IBM code as the reference helped the IBM Token Ring technology to shake itself out much faster than it would have otherwise. The FDDI players worked together to try and understand what had been done by the other players. There was never finger pointing, nor was there grumbling about how people implemented the specs. Write it off as a new technology’s growing pains.

continued on next page

The INTEROP 90 Show Network (*continued*)

25 miles of cable

Now, I am sure you can imagine that it was much more complicated than this. Saying we ran wires here or there does not adequately describe the magnitude of our undertaking. The cabling for INTEROP 90 was delivered in late June 1990, and cable harnesses constructed during the July 4th week when the convention center is traditionally empty. (We call this event the "July Wiring Party.") Originally, we thought we were going to be running something around 12 miles of cable. This was based on the assumption that there were 210 booths, all would get at least 1 drop, and the ceiling height was 40 feet. Well, let me say this, we were *very* conservative in our estimate. We didn't want any unpleasant surprises when we decided to run wire. Fortunately, Peter DeVries, the shownet director, was much more detailed-orientated than we were. He counted *every* single wire. He allowed for us to have some slack, and ordered 25 miles of cable. Now, usually, I over-order by 25%, and while the 25 miles of phone wire included some slack, it is obvious he was much more thorough than we were. In the end, we had only about 3 miles of cable left when we constructed the harnesses in July.

In a perfect world, we would have stopped there, and had a pleasant job simply "rolling it out" in October. Unfortunately, several things happened. First off, our installation time was cut from 12 hours to 8 hours. Second, in that 8 hours, we got stuck adding another 7 miles of cable to the harnesses. (These were "engineering mods" resulting from network requests arriving after the deadline, last-minute demo participants and so on).

As with previous INTEROPs, we ran down to the end trying to get the complete show up and running in time for when the doors opened Wednesday at noon. As the years have worn on, we have assembled a large group of some of the best people in the industry who help us assemble this large, complex network. This year, the building joints that gave us so much trouble in previous years were handled without any intervention by the people trying to get the network together. The terminal rooms, the microwave connection to the other building, and the T1 connection were all up and providing people with Internet connectivity for e-mail access Tuesday morning. Sometimes the connections were not as reliable as they should have been, because we were trying to play with the backbones ourselves while you were trying to use them to read your mail.

Management tools

In general, there are always things you could have had that would have made your life much easier. The things we needed were many and varied. It would have been very useful to have had a "real" routing protocol running on the spine. As it was, we had to use RIP, because this was the only common protocol the routers spoke. At INTEROP 91, we hope to use OSPF, which should provide us with a much more robust and dynamic network. The FDDI problems were complicated by RIP's inability to deal with the complex, and rapidly changing, topology. We found that the easiest way to manage the network was with *ping* and *Telnet*. The SNMP management stations we had were useful, but only when someone was around who was familiar with them to run them for us. Several of the vendors present tried to give us crash courses in how to use their management stations, but, in the heat of battle, they (the management systems) were not as friendly to use as *ping*, or *tracert*. Perhaps INTEROP 91 will have tools that will be more useful for managing a network that is falling down around your feet, as you are trying to bring it up for the first time.

Reflections

Over the years, this show, and the network it provides, have given me the opportunity to work with some of the best people in our industry. People whom I would have just brushed against in the business world, I have now seen with their defenses down, sleep deprived, stumbling around at 4am, asking if they could do anything to help get things working. There were times when I asked myself about why one would get involved in this kind of endeavor. It is not out of any devotion to a company, or any great financial incentive. It is brought about by the desire to contribute to a great monument to the achievements of your profession, for a few moments. Something decadent, like creating one of the most complex networks you will ever have the opportunity to see running, appreciate the scope of it, and then destroy it three days later. There is a certain sick, twisted joy in proving that it can be done on your terms, and, after enjoying it for a brief time, taking it apart, and returning to your normal job.

I suppose this is an attempt to answer the "why" question. As before, if you don't understand it, forget it. I am not really sure I understand it either. For me, and the shownet team, INTEROP is 9 days of a surreal world of high-tech marvels and toys. The let down afterwards is incredible. The rush of doing it is indescribable.

[Ed.: Stev Knowles will be back to build the INTEROP 91 show network in October. Following the trend of the last 4 years, the network is larger and more complex than before. This year, an added complication is a second exhibition hall on the other side of the street—and no, there aren't any tunnels providing cable access across the street, so the connectivity will be implemented by means of a microwave link, and possibly via public T1 circuits].

For further reading

- [1] Jacobsen, O., "INTEROP 88 Conference Report," *ConneXions*, Volume 2, No. 11, November 1988.
- [2] Almquist, P., "The INTEROP 88 Network—behind the scenes," *ConneXions*, Volume 3, No 2, February 1989.
- [3] Dern, D., "Highlights from INTEROP 89," *ConneXions*, Volume 3, No. 11, November 1989.
- [4] Knowles, S., "The INTEROP 89 Network: From one of its builders," *ConneXions*, Volume 4, No. 2, February 1990.
- [5] Dern, D., "INTEROP 90 Wrap-Up: A View from the Floor," *ConneXions*, Volume 4, No. 11, November 1991.

STEV KNOWLES is a hacker at FTP Software. (This referring to the "old school" of hackers, when being a hacker was a title bestowed upon you by others. It is not the only title bestowed on Mr. Knowles by others, but it is the only one we could get past the publisher.) Mr. Knowles has been involved, to one degree or another, in all the Shownets that have been staged at the INTEROP trade shows over the years. The networks don't seem worse for the wear, nor does FTP Software. He also manages to stir up trouble at IETF meetings. He has no degrees, certificates, or warrants outstanding for his arrest. He can be reached by Internet e-mail as stev@ftp.com.

All Your FDDI Questions answered at INTEROP 91

by Mark S. Wolter, National Semiconductor Corporation

FDDI a Reality

INTEROP 91 is nearly at hand and the FDDI community is poised for another great show. FDDI representatives and their products will be involved in all aspects of the conference and exhibition. The cooperation between vendors will once again produce the most comprehensive FDDI show available, proving that FDDI is a viable technology available today.

In-depth tutorial

Because of the wide-spread interest in learning the details of the FDDI protocol and how it is installed into existing networks, INTEROP will be hosting an in-depth tutorial entitled "Introduction to the FDDI Protocol and Its Networking Applications." This tutorial will cover the existing FDDI protocol, explaining the implementation and purpose behind the protocol, and its use in networks. Reliability, management, and performance enhancements provided by the standard will be compared to previous LAN protocols. The capabilities provided by the protocol's required management facilities will be explored. The finer points of the installation of an FDDI cable plant will be covered with reference to typical existing cabling systems. The necessary considerations for the internetworking of FDDI networks and other LANs will be included with a brief definition of bridging and routing functions. This tutorial covers the FDDI protocol in detail, and provides insight into networking equipment and procedures required to install and maintain an FDDI network.

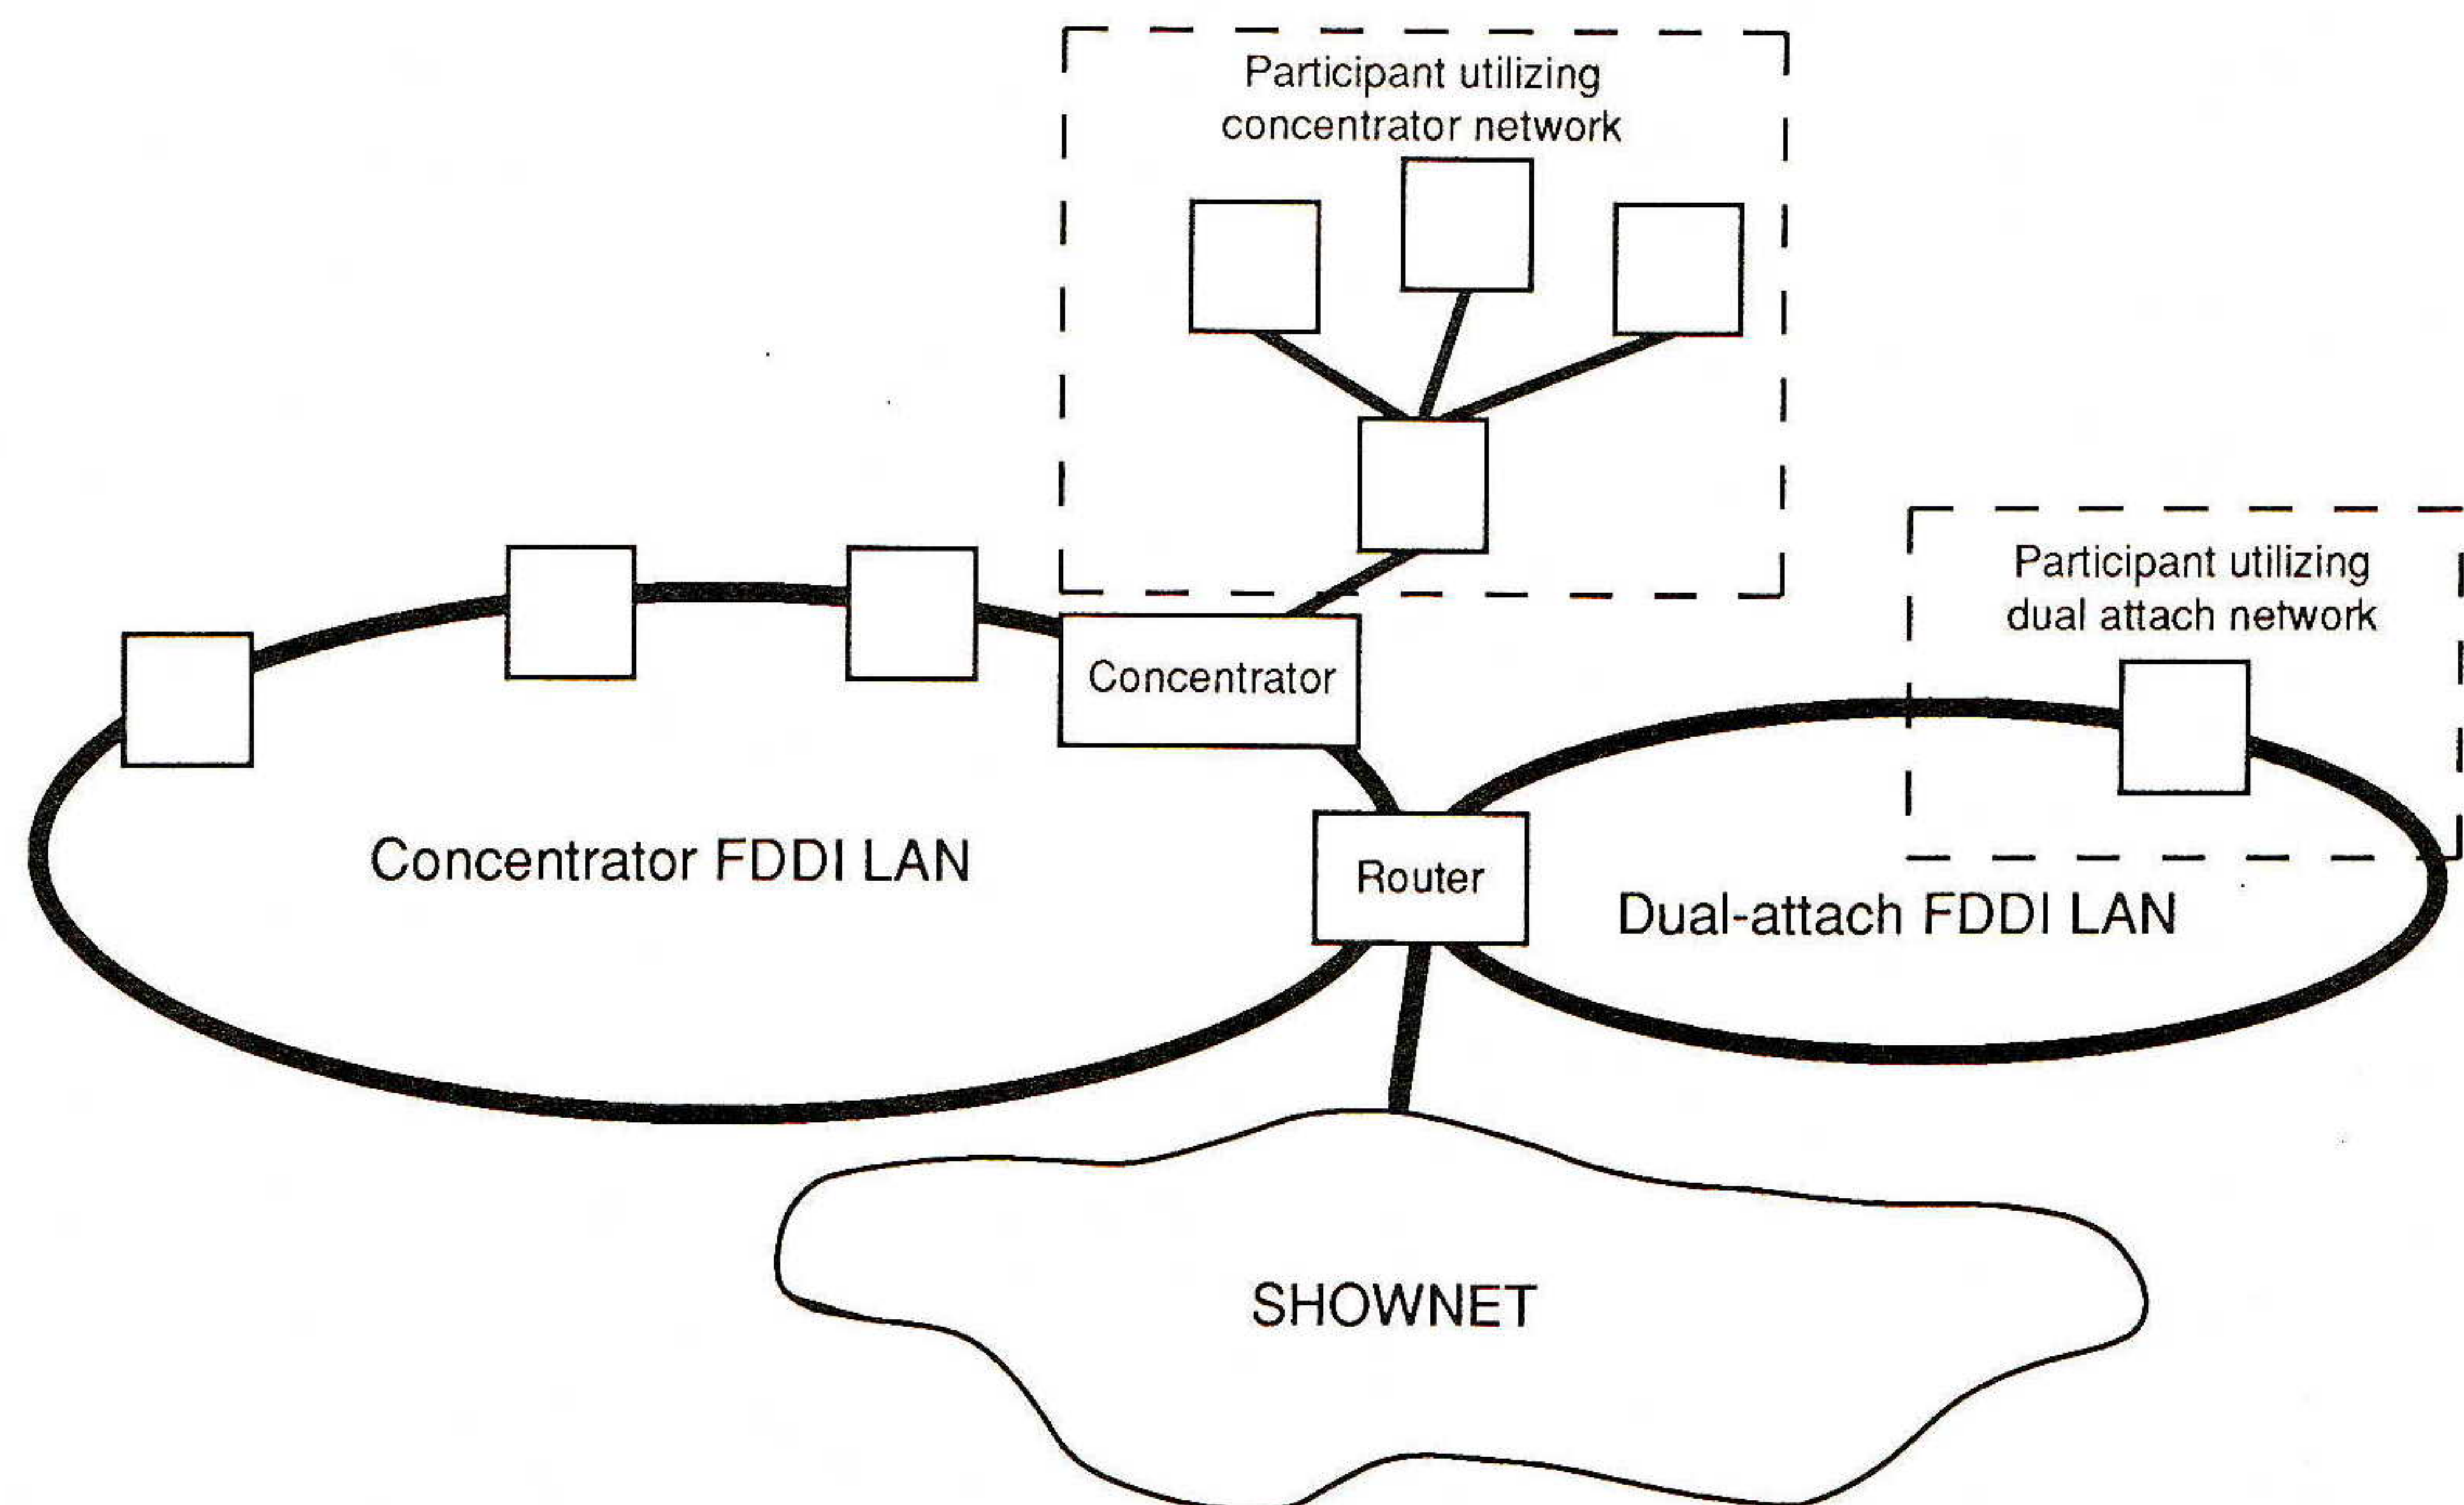


Figure 1: INTEROP 91 FDDI Demonstration Network

Exhibitions

The exhibition will include an FDDI backbone for the Shownet, interconnecting several different networks on one FDDI ring. In this network, FDDI is providing a high speed backbone that can support the traffic generated by the large number of nodes connected to the Shownet. The use of this network will be transparent to the Shownet participants, except for the increased network availability.

An FDDI demonstration network will also be assembled, involving over forty vendors this year. Last year, this network included thirty-seven participating companies and over seven and a half miles of fiber. This year's topology will include two separate FDDI LANs, one offering participants connections through concentrators while the other offers dual-attach connections. The concentrator-based FDDI LAN will be hierarchically organized with connections going to concentrators in seven pedestals located throughout the exhibition hall which will then be connected to the Showcase patch panel. A router connects the two FDDI LANs in the demo Showcase with each other and with the INTEROP Shownet.

The FDDI Showcase booth will focus on FDDI in real-world environments, the theme of this year's FDDI demonstration. Displays addressing the use of an FDDI network will cover topics related to the installation, management, and use of the network. Examples of cabling will be shown along with example topologies for different networking needs. Management facilities will be used to monitor and map the demonstration network. Typical applications made practical by the increased bandwidth of an FDDI LAN will be featured in the booth. Participants from the demonstration will be manning the booth and are available to discuss issues and address questions concerning the demonstration and FDDI in general. Drop by and ask the experts, and don't forget to pick up your *Pocket Guide to FDDI*.

Participants in the demonstration are also cooperating to assemble several applications on the ring that will be displayed in their own booths. Applications available between participants will include video imaging, signal processing technology, network-based services, network file sharing system, and network management and monitoring demonstrations, among others. Participants are also involved in developing separate demonstrations among themselves, but transmitted on the FDDI demonstration network. These applications will provide a realistic load on the network.

Panel discussions

This year, several panel sessions will accompany the exhibitions, and the FDDI demonstration group is well represented here as well. There will be a total of four sessions organized, including panelists from both the customer and vendor communities. Two sessions will focus on the implementation experiences of customers who have designed, installed, and used FDDI networks within their organizations. Attendants will have the opportunity to take advantage of other customers' experiences and ask questions regarding FDDI networks currently in use. Another session will cover the future directions of FDDI regarding SMT, FDDI topologies, and media discussed by several members of the ANSI FDDI X3T9.5 standards committee, while the forth session will address the planning required for the installation of a network and the implementation and use of network management facilities with insightful perspectives provided by FDDI vendors. Attendants are encouraged to participate in the discussions during each of these sessions.

Conference sessions

Conference sessions will cover the current use of FDDI in differing applications and the future changes under development to enable wider use of FDDI. The first session, "FDDI Technology and Architecture" will describe what FDDI is and how it is used, including the use of FDDI in a backbone application and as a front-end network connection providing full bandwidth to desktop workstations.

FDDI Questions answered at INTEROP 91 *(continued)*

The capabilities of SMT to manage an FDDI network, and how SMT is used to manage a station from a remote network will also be covered in this session. Interconnections between FDDI and other LANs by bridges and routers will be covered, as will the design and installation of a practical FDDI network. This session is recommended for conference attendants who want an overview of FDDI technologies.

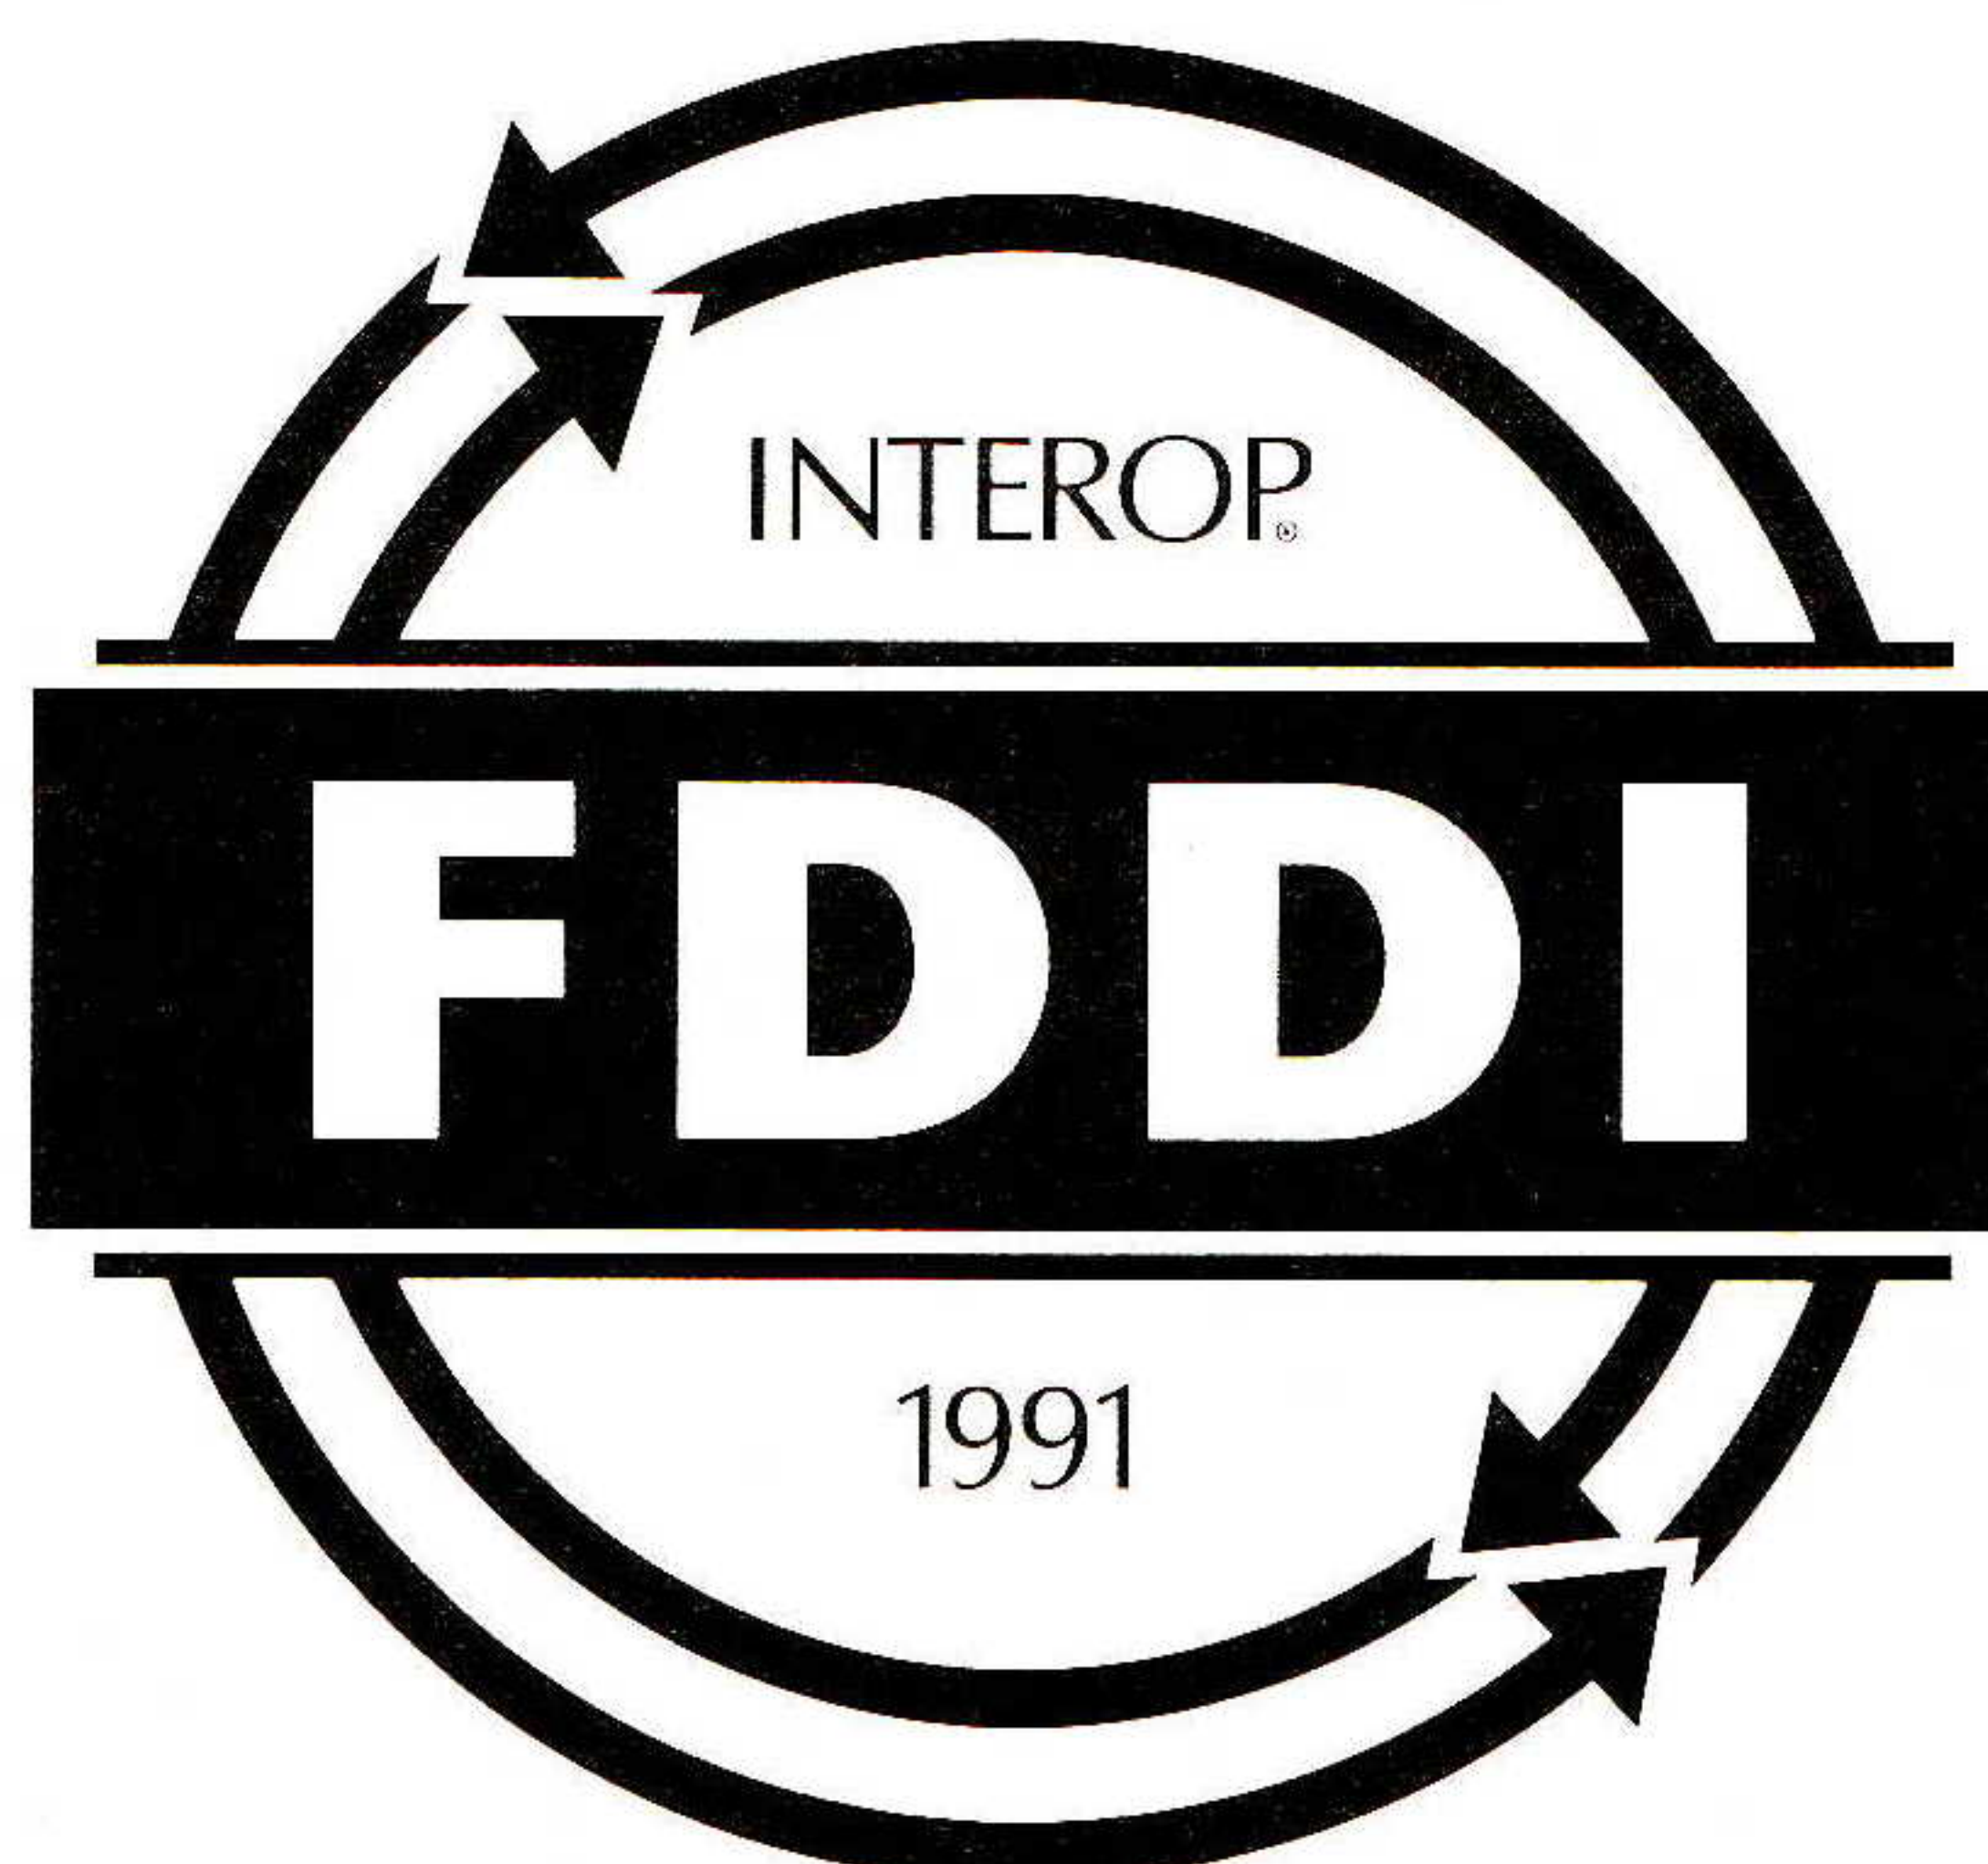
A second session, entitled "New FDDI Technology" will focus on emerging FDDI technologies, and includes additions to the FDDI standards that will be introduced within the next five years. FDDI will be going through several changes, including the reduction of costs of the transmission media and transceivers, the addition of multimedia services, and changes related to the long term trends of LANs in general. Alternative media currently under development, including low cost fiber and twisted pair copper wires solutions will be covered. The integrated services capabilities of FDDI-2 will be discussed. Trends leading changes in the interconnection of LANs across WANs and higher data rate protocols involved in gigabit LANs will be addressed in this session as well. Each presentation will cover practical implementation issues for appropriate future applications.

[Ed. A second FDDI article, entitled "ANSI X3T9.5 Update: Future FDDI Standards," will appear next month].

References

- [1] Wolter, M., "Fiber Distributed Data Interface (FDDI)—A Tutorial," *ConneXions*, Volume 4, No. 10, October 1990.
- [2] "Solutions Showcase Demonstrations for INTEROP 91," *ConneXions*, Volume 5, No. 8, August 1991.
- [3] Wolter, M., "ANSI X3T9.5 Update: Future FDDI Standards," *ConneXions*, Volume 5, No. 10, October 1991.

MARK WOLTER is an applications section head for the Advanced Communications Group at National Semiconductor and participates on the ANSI FDDI committee. He has written papers and presented seminars for several systems conferences. He has experience in the design of both FDDI networking systems and ICs. He can be reached as wolter@berlitz.nsc.com.



Letters to the Editor

Ole,

I'd like to make a correction to the article "Networks: From Technology to Community," which appeared in the July 1991 issue of *ConneXions*. It's not really a big point, but in the section labeled NFS where it mentions X it says, "*The X Window System* was invented by MIT Project Athena..." which is incorrect. If the "invention" is to be attributed to any one group at MIT it would be the Laboratory for Computer Science (LCS) where it was first "invented" by Robert Scheifler, based on earlier work elsewhere. Some of the development on the later versions was contributed by Project Athena, as was a substantial part of the *X Window System Toolkit* (which is not the core). If you need to refer to this in a real short expression, you might try the following which is more correct:

"The X Window System was developed at MIT by the Laboratory for Computer Science (LCS) and Project Athena."

Yours for detailed correctness...

*Michael A. Patton, Network Manager
Laboratory for Computer Science
Massachusetts Institute of Technology*

Ole,

As always, by the time my article ["Another use of the Internet: Libraries Online Catalogs,"] was published (and you had quick turnaround), technology had taken another step forward. Thanks to Brewster Kahle and Thinking Machines Corporation my library guide mentioned in my July 1991 *ConneXions* article is now available through WAIS. Since I know that you are publishing an article on WAIS in the near future, I thought the readership might like to have the WAIS source record on my library guide.

```
(:source
:version 3
:ip-address "129.120.1.42"
:ip-name "sol.acs.unt.edu"
:tcp-port 210
:database-name "library"
:cost 0.00
:cost-unit :free
:maintainer "billy@unt.edu"
:description "Server created with WAIS
release 8 b1 on Mon Jul 8 08:46:21 1991
by root@sol"
)
```

*Billy Barron
VAX/Unix Systems Manager
University of North Texas*

[Ed.: Indeed, the article "An Information System for Corporate Users: Wide Area Information Servers," by Brewster Kahle and Art Medlar will appear as soon as the INTEROP 91 crush quiets down].

Call for Presentations/Papers

MULTIMEDIA '92—The 4th IEEE COMSOC International Workshop on Multimedia Communications will be held April 1–4, 1992 in Monterey, California.

MULTIMEDIA '92 is the 4th international workshop dealing with multimedia communication systems and services as well as the creation, processing, storage, transmission and human utilization of multimedia information that gives a strong impact on realizing a future "Information Society."

MULTIMEDIA is a forum for telecommunication networks, terminals and applications specialists to meet, discuss and understand each other's objectives, requirements and constraints. MULTIMEDIA '92 will be more specifically devoted to applications in the area of computer and human communication.

The workshop will be held at the Hyatt Regency in Monterey, California. The participants will be limited to 150.

Topics

Topics for MULTIMEDIA '92 include:

- Network Architecture and Implementation
- Node Architecture and Implementation
- Protocol Design, Analysis and Engineering
- Human Interface
- Multimedia Terminals & Databases
- Standardization Activity & Experience
- Relevant Applications
- Network-Supported Collaboration

Submissions

Authors are invited to submit an extended abstract of 1500 words in English, clearly indicating a contribution in the areas of multimedia communications shown above. Any other topics or related subjects are welcome. Authors should submit three sets of abstract with their complete address (including Fax & E-mail) to:

J. Joaquin Garcia-Luna
 Technical Program Chair,
 MULTIMEDIA '92
 SRI International
 333 Ravenswood Avenue
 Menlo Park, CA 94025
 Phone: 415-859-5647
 Fax: 415-859-6028
 E-mail: garcia@sri.com

Important dates

Abstract Submissions due: October 1, 1991

Notification of Acceptance: December 1, 1991

Call For Papers

Papers are solicited for the *Silicon Valley Networking Conference* (SVNC-92) to be held April 27-29, 1992 at the Santa Clara Convention Center in California.

Topics

Papers are solicited in the following areas:

- Distributed Systems
- Network Management
- Advanced File servers
- Standards activities
- Internetworking
- The X Window System
- High Speed Networking
- Wireless Networking

SVNC typically attracts over 400 engineers every year and is a nice forum to discuss system design architecture and solutions to complex networking problems.

Submissions

If you are interested in presenting a paper, please send an abstract of the paper before September 23, 1991. If accepted for submission, a rough draft of the paper should be submitted before December 15, 1991 and camera ready copy should be submitted before January 25, 1992. Please include your Address, telephone number and fax number in the abstract, and send it to:

B.V. Jagadeesh
3Com Corporation
5400, Bay Front Plaza
Santa Clara, CA 95052
408-764-5169 • bvj@3Com.com

Exhibits

SVNC also offers vendors to exhibit their products on a table top during the conference. Companies interested in demonstrating their products please contact the above through e-mail or telephone.

ISDN Mailing List

ISDN@List.Prime.COM is an electronic mailing list which has been in existence for some time. The list is for the discussion of the technical aspects of ISDN. There are presently 1400+ subscribers in more than 25 countries, including many of the (would be) ISDN providers.

The ISDN list is not moderated. All submissions are automatically reflected to the entire list of subscribers. Note that discussion of voice telephony may be more appropriately carried on in the TELECOM Digest and the USENET newsgroup (comp.dcom.telecom). This is, of course, at the contributor's discretion.

You can receive this list by SMTP/X.25/ISDN; see RFC 1090 and ask ISDN-Request@List.Prime.COM for more information. An archive of sorts is available via anonymous FTP, on host tiger1.prime.com in file [pub/isdn/](ftp://tiger1.prime.com/pub/isdn/)*

To add your name to the list, send to a message to:

ISDN-Subscribe@List.Prime.COM

Submission to the list should be sent to:

ISDN@List.Prime.COM

The list coordinator is:

Robert Ullmann, 508-620-2800 x1736 • Ariel@Relay.Prime.COM

Call for Papers

The *IFIP International Conference on Upper Layer Protocols, Architectures and Applications* will be held May 25–29, 1992. The conference is sponsored by IFIP Working Group 6.5 and hosted by the University of British Columbia, Vancouver, Canada and OSIWARE Inc., Vancouver, Canada. [Ed.: Note new dates!]

Conference outline

The purpose of the conference is to provide an international forum for the exchange of information on the technical, economic, and social impacts and experiences with upper layer protocols, architectures and distributed applications. The conference format will be two and a half days of conference paper presentations combined with one half a day of workshops.

Topics

Papers are desired in—but not restricted to—the following topic areas:

- Design, implementation, experience with distributed applications
- Application layer user agent models and designs
- Application layer architectures
- Application layer programming environments
- Application communication protocols e.g., RPC, RO, RTS and multicast
- Group communication models and services
- Multimedia applications and communications
- Interconnection of upper layer and application entities
- Upper layer conformance and interoperability testing activities
- Security and privacy
- Management and operation of distributed services
- Mobile communications and the application layer protocols
- Upper layer network management and naming
- Presentation and session layer issues
- Abstract syntax notations and transfer syntaxes
- The impact of very high-speed networking on the upper layer protocols

Instructions for authors

Prospective authors are invited to submit for review, unpublished original contributions (not exceeding 5000 words) which describe recent research results or developments on any design or service aspect of upper layer protocols, architectures or distributed applications. Papers that are accepted will be published by North-Holland. A preprint of the proceedings will be provided to attendees.

Deadlines

- | | |
|-------------------|---|
| Today: | Send a message, letter or phone any of the contacts below stating your intention to submit a paper, or stating your general interest in the conference. |
| December 1, 1991: | Full version of papers due for review. |
| March 1, 1992: | Notification of acceptance/rejection. |
| April 15, 1992: | Camera-ready papers required for publication. |

Submissions

Please submit five copies of your paper to:

Dr. Gerald Neufeld
Department of Computer Science
University of British Columbia
Vancouver, B.C, Canada, V6T 1W5
Telephone: +1 604 228 4806
Facsimile: +1 604 228 5485
Internet: neufeld@cs.ubc.ca
X.400: S=Neufeld; OU=CS; O=UBC; P=cdn;
A=telecom.canada; C=ca

or to:

Prof. Dr. Bernhard Plattner
Laboratory of Computer Engineering and Networks
Swiss Federal Institute of Technology (ETH)
ETH-Zentrum CH-8092
Zürich, Switzerland
Telephone: +41 1 254 7000
Facsimile: +41 1 262 3973
Internet: plattner@komsys.tik.ethz.ch
X.400: S=Plattner; OU=komsys; OU=tik; O=ethz;
P=switch; A=arcom; C=ch

Tutorials

On May 25 and 26, 1992, time and space is reserved for a number of tutorials. The following topics are being considered:

- Upper Layer/Application Layer Architecture
- Security Models, Mechanisms and Systems
- Message Handling X.400 (1992)
- Directory Services X.500 (1992)
- Network Management
- ASN.1
- Coexistence and Transition to OSI Applications
- Distributed Application Programming Environments

Write to *ConneXions*!

Have a question about your subscription? Suggestions for topics? Want to write an article? A letter to the Editor? Have a question for an author? Want to enquire about back issues? (there are now more than fifty to choose from; ask for our free 1987–1991 index sheets). We want to hear from you. Simply write, call, fax, or e-mail to:

ConneXions—The Interoperability Report
480 San Antonio Road, Suite 100
Mountain View, CA 94040-1219
USA
Phone: +1 415-941-3399 or 1-800-INTEROP (Toll-free in the USA)
Fax: +1 415-949-1779
E-mail: connexions@interop.com

CONNE^XIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

FIRST CLASS MAIL
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

ADDRESS CORRECTION
REQUESTED

CONNE^XIONS

EDITOR and PUBLISHER Ole J. Jacobsen

EDITORIAL ADVISORY BOARD Dr. Vinton G. Cerf, Vice President,
Corporation for National Research Initiatives

A. Lyman Chapin, Chief Network Architect,
BBN Communications

Dr. David D. Clark, Senior Research Scientist,
Massachusetts Institute of Technology

Dr. David L. Mills, Professor,
University of Delaware

Dr. Jonathan B. Postel, Communications Division Director,
University of Southern California, Information Sciences Institute

Subscribe to CONNE^XIONS

U.S./Canada ☐ \$150. for 12 issues/year ☐ \$270. for 24 issues/two years ☐ \$360. for 36 issues/three years

International \$ 50. additional per year (Please apply to all of the above.)

Name _____ Title _____

Company _____

Address _____

City _____ State _____ Zip _____

Country _____ Telephone () _____

☐ Check enclosed (in U.S. dollars made payable to CONNE^XIONS).

☐ Visa ☐ MasterCard ☐ American Express ☐ Diners Club Card # _____ Exp. Date _____

Signature _____

Please return this application with payment to:

CONNE^XIONS

Back issues available upon request \$15./each
Volume discounts available upon request

480 San Antonio Road, Suite 100
Mountain View, CA 94040 U.S.A.
415-941-3399 FAX: 415-949-1779

CONNE^XIONS